

EXHIBIT A

CONFIDENTIAL

REDACTED

Secure Data Interchange

Field of Invention

- 1 This invention relates to systems for the personalization of information delivery, including the delivery of
2 advertisements, product information, news and features. The system of Secure Data Interchange provides users and
3 vendors with absolute control over profile information, while enabling focused targeting of information, and profile
4 interchange between different entities. The system provides the technical infrastructure for a market for profiles,
5 evaluation, and personalized information and product delivery.

6 Problem

- 7 The introduction of cheap and powerful new information technology allows manufacturers, service providers, and
8 stores (on-line and off-line) to collect information about the transactions and preferences of customers and users
9 cheaply and efficiently. Moreover, new network connectivity enables different vendors to exchange profiles for
10 common customers, either statically or dynamically, in order to build broad and detailed profiles across vendor
11 domains. There exist many potentially powerful synergies between the data sets that are collected by different vendors
12 and service providers, that can be leveraged to provide appropriate services and products to customers. When analyzed
13 with the proper statistical tools these data sets can reveal fundamental patterns in the behavior of users, and enable a
14 vendor to provide appropriate information to a user. Furthermore, access to user-profiles collected by other vendors can
15 enable vendors to provide focused information delivery to first-time users, and also cross-market services with other
16 appropriate vendors.

- 17 Electronic intermediaries that monitor the activities of users across different vendors and service providers can also
18 collect data about the products and services that vendors provide. This data can be used, with appropriate analysis, to
19 provide users with advice about relevant services and products. User profiles can be used to identify the goals,
20 preferences and interests of users, vendor profiles can be used to relate the services and products provided by vendors
21 to the profiles of relevant users. Users can benefit because they can find information more readily, and vendors can
22 benefit because they can reach potential customers more easily.

- 23 The problem with the ease with which data can be collected, and the ability to readily integrate information that is
24 collected in diverse transactions and activities, is that all of this information represents a significant challenge to the
25 privacy of individuals. A visitor to a web site does not even have to buy anything for information about his activities to
26 be monitored, recorded, and passed onto other web sites in future interactions - for example through the use of
27 "Cookie" technology that is largely transparent to end users.

- 28 Thus, although there are many advantages to building large intra-and cross-industry databases, most companies by
29 necessity must keep their data to themselves, and individual users must be on their guard for unrequested and
30 inappropriate solicitations from vendors who misuse their personal information. What is required is a system that
31 enables data exchange and analysis within a secure framework that ensures privacy and protects against misuse of
32 personal and transactional information. There is a conflict between the privacy rights of customers and effective
33 marketing, with a focus on using information gathered about customers to refine offers to users.

- 34 Both vendors and users can benefit from the exchange of some information on transactions and personal preferences,
35 from users to vendors, and between vendors. In the same way that a large amount of information leads to information
36 overloading, and makes it difficult for an individual to find the information that he/she is searching for-a dynamic
37 marketplace with many vendors, products and services can make it difficult and expensive for a user to locate an
38 appropriate product or service. Similarly, vendors would like to offer their services and products to users who are most
39 likely to benefit, and most likely to make purchases. There would be clear and powerful synergies if technology existed
40 to enable:

- 41 (a) secure evaluation of the value of data to a requesting party;
42 (b) secure data transfer in an environment that ensures that data privacy rules are protected at all times.

BEST AVAILABLE COPY

CONFIDENTIAL

1 Vendors would find benefit in sharing data with other vendors; this would deepen their understanding of their
2 'customers' behaviors and preferences, especially if certain customers were traceable across several data sets. One could
3 imagine an on-line music store sharing data with an on-line ticket dealer. Firstly, they would be able to augment their
4 mailing lists with each other's customers. Secondly, they would increase their understanding of what kinds of music
5 particular customers prefer. For example, by analyzing the particular types of music that individual customers prefer,
6 the ticket vendor could target appropriate concerts to each customer.

7 Users (individuals) would benefit from sharing data with other users (individuals). This is already evident in the
8 popularity of news groups and web pages catering to individuals with shared interests. By learning what other people
9 with similar tastes and preferences have discovered and enjoyed, a user can sidestep information overload in the search
10 for personally satisfying information.

11 Vendors and users can benefit from each other. An obvious example would be in the use of collaborative filtering for
12 the marketing of targeted promotions; rather than being deluged with coupons and advertisements that are of absolutely
13 no interest, a user would benefit by being presented with advertising that is highly relevant. In the process, the vendor
14 would increase advertising response rates, boosting overall efficiency. Users could also benefit from the personalization
15 of content at vendors' web pages, and well focused banner advertisements at other web sites that they visit.

16 While the above scenarios demonstrate the potential benefits of the sharing of data across different parties, they also
17 present the possibility of the misuse of data: vendors could sell each other's private data to rivals, and users'
18 information could be used against them. At present there is no technological solution that provides the many benefits of
19 information exchange, as outlined above, but without the ability to misuse data.

20 This invention relates to a new technique for encoded storage and communication of information that allows aggregate
21 information to be recovered while protecting the privacy of information pertaining to any one individual. This
22 invention also relates to, but is not limited to, a particular system incorporating the above-mentioned technique in
23 generating target profiles in a system for customized electronic identification of desirable objects. In one embodiment
24 of this system a profile is maintained for each user that, in the general case, records both the user's demographic
25 attributes and a record of the user's buying habits or other expressed preferences. However, the profile is maintained in
26 an encoded form to be described. With this encoding it will be impossible to extract particular information pertaining to
27 an individual user. It will only be possible to get information about large groups to which the user belongs.

28 In this system a user's profile is maintained by continuously acquiring information about past decisions by the user and
29 other "similar" users. By encoding the feedback from the user using the new technique the particular decisions made by
30 the user are not revealed to anyone else.

31 Information gathered about commercial or other activities by individual users can be used to provide a higher quality of
32 service to such users in a number of applications such as electronic commerce and governmental services. Competitors
33 and law enforcement agencies can also use such information. Thus it is important to allow access to such information
34 only as permitted by the user or as permissible by law.

35 Solution

36 The above-described problems are solved, and a technical advance achieved, by the system of Secure Data Interchange.
37 The Secure Data Interchange architecture and system enables the profiles of users across many vendor types to be
38 combined into portfolios and system-level aggregate information to the extent that a user specifies, within a system that
39 is secure to manipulation from vendors. SDI presents a technical solution that alleviates some of the tensions and
40 conflicts that exist between privacy and focused information delivery. SDI allows vendors to personalize, while users
41 can remain in absolute control of their personal profiling information.

42 Profile information is released to a central Secure Data Interchange data warehouse, or to data warehouses provided
43 operated by vendors or third parties, and can be both anonymized and randomized to protect the privacy of individuals,
44 while allowing detailed statistical analysis and model building for dynamic on-the-fly personalization of services.
45 Secure Data Interchange resolves many of the conflicts that exist between the benefits of personalization and well-
46 focused products and services, and protecting the privacy of users. In particular, privacy is protected through
47 restrictions on the amount of explicit information that is released about users, controls over the mechanisms with which
48 vendors can track users in between sessions, and protecting the release of implicit data (e.g. "clickstream" data, that is
49 not transactional).

CONFIDENTIAL

1 The primary application of SDI is to a Internet-based electronic commerce system, with individual customers "users"
2 and larger vendors connected over a network of clients and servers. As such, SDI can be implemented as described
3 with existing standards and protocols, including (but not limited to) the HyperText Transfer Protocol (HTTP) for client-
4 server communication, extensible Markup Language (XML) for embedding meta data in Web documents, Java applets
5 and Java script for client-side processing, and encryption methods, such as Secure Sockets Layer (SSL) and the X.509
6 standard for secure information transfer. Secure Data Interchange also complements and significantly extends current
7 proposals for open profiling (Open Profiling Standard OPS), and user privacy controls (Platform for Privacy and
8 Preferences), that are active projects of the World Wide Web consortium. The key advance provided within SDI is the
9 ability to manage privacy and control personalization within an integrated system.

10 SDI allows users to specify privacy and data-release policies, and control the aggregation of information across
11 different vendors. Furthermore SDI enables vendors to augment detailed transaction-based information with broader
12 information about users that is collected from their extended interactions with other information and service providers,
13 again within a privacy protected system. SDI provides a user with complete control over her identity as she browses
14 the Internet and makes on-line purchases, placing the user in absolute control over what information each vendor can
15 collect or receive about the user.

16 SDI can manage the information that a user wishes to release to vendors when he/she registers with a new system. A
17 client-side SDI enabled proxy (or browser plug-in) controls information revelation, and provides support for
18 pseudonymous and anonymous interactions. SDI has a distributed architecture, with accurate and complete profiling
19 information about a user maintained on trusted clients, and randomized (and potentially anonymized) information
20 pushed to centralized data warehouses, information that is still valuable for trend analysis and model-building.

21 The main technical solutions that are used within the system of SDI to enable privacy guarantees and personalization
22 are: (1) Information-theoretic tools, such as releasing randomized profile information to third-parties and vendors, and
23 removing identifying information from messages between users and vendors; (2) Distributed secure processing, such as
24 locating user profile information on the client machine of the user, and using local secure processing to personalize
25 generic information provided by vendors; (3) Cryptographic techniques for pseudonym generation and validation, and
26 user authentication; (4) multilevel collaborative filtering techniques.

27 The Secure Data Interchange provides a number of primary functions. First, the system provides a secure environment
28 where data can be collected, aggregated across a number of vendors and users, and analyzed such that the privacy and
29 usage requirements of the user and vendor are protected. We are able to use technological solutions to provide
30 guarantees about the information that is made available to vendors. Second, the system manages the privacy policy of a
31 user, through the control of information and automatic management of pseudonyms. Third, the system allows vendors
32 to provide personalization to users without accessing potentially sensitive and valuable profile information about users.

33 The system also provides support for electronic commerce functionality, such as providing a mailing service for
34 targeted and authorized solicitations to a virtual and anonymous mailing list, and providing support for a system of
35 certification and anonymous payments. The Interchange can also provide a privacy-protected market place for customer
36 information, where certain types of data can be purchased, rented or sold by vendors in order to enable synergies to be
37 realized between the data sets of independent service providers and vendors.

38 1. Definitions

39 For the purposes of this patent it is useful to separate the users of the Secure Data Interchange into two general classes:
40 users and vendors.

41 A user is an individual (person), or possibly a group of individuals (people) that share similar interests, that transacts
42 with multiple vendors, information providers, and service providers, and has an interest both in receiving personalized
43 service and protecting his/her privacy with respect to those transactions. In particular, a user might require that two
44 vendors with which he has had transactions cannot build an integrated profile on the basis of the user's independent
45 transactions, with the goal of using information about one user-vendor session to provide a refined service in another
46 user-vendor session. As another example, users might state particular types of vendors that are able to exchange
47 information about the user, but for restricted purposes. A user is connected to the Internet through a dedicated host
48 machine running client software, such as an Internet browser with an SDI plug-in or secure Java applet. Users include
49 individuals shopping and browsing on the World Wide Web, whose extended purchasing and interaction behaviors can
50 allow vendors to target their products and services.

CONFIDENTIAL

1 A vendor may be an information provider, or a provider of goods and/or services, connected to users and SDI through a
2 dedicated server and an SDI proxy server. Unlike users, vendors have less interest in pseudonymous interactions with
3 different users, because vendors will typically have an interest in developing a strong identity. Like users, vendors may
4 wish to exchange information on transactions with other vendors while maintaining tight control over the type of
5 vendors that receive information, and also over the exact use that is made of the information. For example, a vendor
6 might be willing to sell information on its transactions to another vendor that is not a direct competitor, but able to
7 cross-sell related products. A vendor might also like to purchase information from users about the transactions that they
8 have engaged in with competitors. Vendors include stores, such as grocery stores, car dealerships, on-line music stores,
9 on-line bookstores, and also information providers, such as web portals and on-line newspapers. Vendors also include
10 external organizations, such as credit agencies, HMOs, schools and police stations, who may hold information on
11 individuals.

12 2. Architectural Overview

13 In this section we provide an overview of the invention, describing both its top-level architecture, and describing the
14 key technologies that we use to support dynamic and powerful data synergies within a framework that explicitly
15 protects the privacy of users. Figure 1 shows the top-level architecture.

16 The system of Secure Data Interchange is a technical solution to the problem presented by the conflicting goals of
17 providing personalized information/products/solicitations to users (beneficial to users and vendors) and the
18 rights/desires of individuals to privacy in transactions. In particular, a user might desire to control the vendors that have
19 a lot of information about its preferences, and prevent certain vendors/entities from gaining information about
20 particular transactions. A user might also like to be able to "keep control" of information that is potentially valuable to
21 vendors and advertisers, and allow only restricted access to that information — and extract a payment for such access.
22 This is the user-centric model of SDI. Similarly, vendors would like to be able to exchange profile/market-research
23 type information with other vendors, but not competitors, and require a secure system for protecting information that
24 they make available.

25 The main data structure within SDI is the profile. A profile is a record of pertinent information about a user, for
26 example demographic information, information about web pages browsed, content information from web pages
27 browsed, recent on-line transactions, etc. Within SDI we associate profiles with an OWNER, and sign them with a
28 ALLOWED-USE key to indicate what the owner has restricted the use of profile information to. The main technique
29 used to support privacy within SDI is a pseudonym, that can fully represent a user in all transactions without providing
30 any information about a user's true identity.

31 Profiles may be located on client-machines, with users; at intermediate level SDI servers, for example at ISP servers; at
32 vendor-level SDI servers; and also in a central SDI server. Users define privacy policies on their client machine, and
33 the client machine ensures that all user-vendor interactions are consistent with a user's policy, and that they cannot be
34 abused — even by a malicious vendor. Client-level proxies can monitor users can build profiles (for example, as users
35 browse and interact with vendors), and vendor-level proxies can monitor users that interact with the same pseudonym
36 in multiple sessions. Both client-level proxies and vendor-level proxies can provide profiles to the central SDI server,
37 with associated conditions of use. The central SDI server can aggregate and validate new profiling information that is
38 received. Client-level and vendor-level proxies can periodically request profile updates.

39 2.1 Main Modules

40 Users interact with the system of Secure Data Interchange through clients, that are general purpose computers with
41 memory and a connection to a network of other computers (clients and servers), for example via the Internet. The User
42 represents an individual that is interested in performing multiple on-line and off-line transactions, with multiple
43 vendors and other users, within the managed data interchange and privacy framework of SDI. Each user physically
44 interacts with the Internet via his/her host machine. The host machine need not be the same machine for all sessions,
45 and could include a user's home PC, work PC, and also any portable devices that the user has configured for SDI. The
46 host machine runs client software, which will include an Internet browser, and also the client-level SDI proxy server, or
47 appropriate browser plug-ins to make the browser SDI-enabled. A host machine is in general a computer with a CPU,

CONFIDENTIAL

1 disk storage, a network connection, and main memory. We assume that the user trusts the host machine and client
2 software.

3 Vendors interact with the system of Secure Data Interchange through servers, that are general purpose computers with
4 memory and a connection to a network of other computers (clients and servers), for example via the Internet.

5 In addition to clients and servers, there are SDI modules that run on the machines, termed "client-level proxies" and
6 "vendor-level proxies". The client-level proxy is a core component that interprets user messages and makes sure that all
7 interactions with vendors satisfy a user's privacy and data-use policies. The vendor-level proxy enables vendors to
8 interact with other key SDI entities: client-level proxies and the central SDI database.

9 Another SDI module resides on a gateway between user's and the Internet, to protect the pseudonymity of a user
10 further. For example, when the user's client machine resides on an intranet (eg that of an ISP) then there is an ISP-level
11 proxy server at the gateway to the Internet that ensures no identifying information is provided to servers. The ISP-level
12 proxy servers also support pseudonymous email addresses for users, and may maintain a database of profiles for user
13 pseudonyms in their user-base.

14 The other main SDI module is a central SDI server, that maintains detailed records of profiles for every user's
15 pseudonym. Having all the information available for access at a unified level allows extensive data mining and
16 collaborative filtering techniques to be applied, but still without violating user privacy and data use policies. In fact, the
17 profile information can be physically distributed, for example located on the SDI-level proxies for each user base.

18 For example, it is possible to provide recommendations to a user under one of his/her pseudonyms through
19 collaborative filtering techniques, even without any other information about the user—just from similar profiles from
20 other users. Similarly, it is possible to recommend new prospects to vendors from the user base of other vendors,
21 without actually providing vendors with profile information about prospective customers. Furthermore, the system of
22 SDI can ensure that vendors can only provide impressions to users if it is allowed within a user's privacy policy.

23 The function of the central SDI server is to make as much use of profile information that is collected (and authorized)
24 by a user for a particular pseudonym, by analysis of the profile information. The goal is not to try to augment the
25 profile with more data—for example by adding demographic information or purchasing transaction records from other
26 vendors. This is impossible within SDI because pseudonyms cannot be associated with a real-world identity unless
27 authorized by a user, and information provided in the profile for a pseudonym is carefully filtered to prevent user
28 identification.

2.2 Overview of System

30 The client-level proxy manages all interactions that a user has with other users and vendors. Essentially, the proxy
31 interprets a user's messages to other users and vendors, and makes sure that a user's privacy and data-user policies are
32 followed. The proxy also maintains up-to-date profiles for users, and allows vendors to personalize information that
33 they provide to users on the client machine, without receiving access to a user's profile. The proxy is also authorized to
34 automate authentication to vendor servers, and release of certain types of information.

35 A key mechanism used within Secure Data Interchange is that of pseudonymous interactions. This allows users to
36 maintain long-term relationships with vendors, and release personal information to vendors, without any other party
37 being able to use the information—is a system of "pseudonyms". Essentially, pseudonyms allow a user to separate its
38 real-life identity from its identity with another user, or a vendor. The client-level proxy is also careful not to provide a
39 vendor with any information that would compromise a user's identity. Pseudonyms provide a very useful middle-
40 ground between total anonymity and complete disclosure. Vendors can still keep useful records of their transactions
41 with a single user, because the system maintains a persistent pseudonym for each user with the same vendor.

42 Users can however choose to interact with each vendor in the system under a different pseudonym, to protect their
43 identity and prevent information transfer across vendors. The control of pseudonyms provides users with a method to
44 control the exchange of profile information between vendors.

45 A pseudonym allows a user to maintain a number of persistent relationships with different vendors or groups of
46 vendors, with complete assurance that the vendors cannot use the pseudonyms themselves to build a profile of the user
47 from the user's sessions across different vendors. A pseudonym provides: Identification, Authentication, Encryption,
48 and Contact information. Pseudonyms are triples: (Pseudonym ID, Private Key, Pseudonym e-mail address).

CONFIDENTIAL

1 Personal information collected through transactions across multiple vendors by the same user, under different
2 pseudonyms is protected from transfer between vendors. Even when vendors themselves maintain local records of the
3 transactions that they have performed with users, and response to queries that they have sent to users, there is no way of
4 providing information in a form that another vendor can combine with its own information and use to personalize its
5 service to the same user. This is because no two pseudonyms can be linked except at the (trusted) client machine. In
6 addition, "randomized aggregates" are used when providing information to a vendor that might allow two vendors to
7 link their pseudonyms. For example, the zip code, age, hair color, and profession of a user might all be useful
8 information for a vendor—but providing all that information accurately to more than one vendor could allow the
9 identity of a user to be compromised and vendors to exchange information.

10 ** DP comment we are also careful to support different keys for interactions between users and vendors -- c.f. new
11 "light security protocol paper".

12 Pseudonymous profiles are still useful in the aggregate, as part of a collaborative filtering system, even when noise has
13 been added to some fields. The system of SDI has a centralized database of profiles and pseudonyms, that can be used
14 for data mining and other collaborative filtering techniques. For example, suppose a user has a pseudonym that he/she
15 uses to interact with all on-line book stores. Then the centralized database can perform collaborative filtering using the
16 profile associated with that pseudonym and other profiles in the database to make personalized recommendations, for
17 example on behalf a vendor that pays to receive such a service.

18 A vendor that belongs to SDI can use the profile associated with a user's pseudonym to provide a personalized
19 interaction with the user. In one version the profile is released to a vendor, and the vendor can push personalized
20 information to the user. In another version, the vendor can push generic information to the user's client, for
21 personalization on the client. It remains important that the vendor does not learn the true identity of a user if the user is
22 to prevent a vendor from providing information on the user to other vendors that know the true identity of the user. The
23 user's personal information is only safe when it is not possible to associate information with anything except for the
24 pseudonym of a user, and only one vendor can interact with the pseudonym (through SDI).

25 We also allow the release of anonymous/pseudonymous profiles to vendors to allow a vendor to provide "in-house"
26 collaborative filtering, without the ability to target any of the users. The pseudonyms are only useful within SDI.
27 Furthermore, because a user's communication channels with vendors is controlled within SDI, we can provide vendors
28 with "rights" to solicit users, and rented (non-transferable) solicitation rights.

29 In a vendor-centric model of SDI we allow vendors to provide information (e.g. profiles about their user-base) to other
30 vendors within a framework that prevents competitors from gaining a competitive advantage. This is possible because
31 one role of the central SDI database is to act as a clearinghouse for profile information, that matches vendors and
32 suggests data synergies. The central SDI database will not release profile information to competitors, according to rules
33 that are designated by the submitting vendor.

34 Furthermore, even when a user designates that a group of vendors can exchange information on the basis of user-
35 vendor transactions, one variation provides each vendor with a unique pseudonym. The only way that vendors can take
36 advantage of profile information from user interactions with other vendors is via the central SDI database, that has a
37 link between the pseudonyms that a user has with each vendor.

38 The system of SDI can track a user's browsing behavior, by cooperation between the client and the ISP-level proxy.
39 The client-level proxy releases sequences of URLs that a user browsers, associated with the pseudonym under-which a
40 user surfs at any particular time.

41 2.3 Ancillary Systems

42 SDI provides ancillary systems that are necessary to support a pseudonymous e-commerce system, for example
43 independent certifying authorities, that are able to validate "once-in-a-lifetime" pseudonyms, that allow a user to prove
44 that it only has one pseudonym for all interactions with a vendor. Certifying authorities can also issue certify properties
45 of the user that owns a pseudonym, for example his/her credit worthiness, his/her age, his/her nationality—without
46 requiring that the user identify themselves to a vendor or another user, and without breaking a link between pseudonym
47 and real-world ID (even with the Certifying authority itself). The issued patent <<Reference here>> includes a
48 description of anonymous payment mechanisms and a secure certificate management system, that we incorporate in

CONFIDENTIAL

1 this patent by reference. Users can gain new credentials through recognized certifying authorities under one
2 pseudonym, and transfer the certificate to another pseudonym when transacting with another server.

3 Similarly, SDI supports a module that allows e-mail to be sent to pseudonyms.

4 SDI also places personal information (eg mailing address of user, credit card number, ...) etc with a trusted third party
5 that has an agent-relationship with the user. When a vendor wants to (for example mail a physical item to a user) it
6 receives certification from the TTP, and then provides that to another TTP with an agent-relationship with the user that
7 mails the item to the mailing address, without the vendor itself ever receiving information about the mailing address.

8 Legacy/demographic information is integrated with pseudonymous profiles, even when SDI has no way to identify a
9 user's real-life identity with a profile, through client-side updates. SDI provides clients with new information, and
10 clients update profile information for pseudonyms with randomized versions of the new information as they choose.

11 Finally, within SDI there is a method to allow users to receive compensation, in the form of rebates and electronic cash,
12 in return from revealing personal information to vendors (even information that is slightly randomized and cannot be
13 shared with other vendors).

14 The technique of blinded signatures is used extensively within our system to reduce the amount of information that
15 third parties, and SDI, have about users. For example, the technique allows a user to create public/private key pairs
16 without giving a certifying party knowledge about the key pair generated. No party within SDI is able to build a dossier
17 that links the pseudonyms of users without explicit information provided by the user. The user has an absolute method
18 to prevent vendors from generating combined user profiles. Blinded signatures are described by D Chaum (??), and
19 incorporated here by reference. Originally developed for the purposes of anonymous digital cash, they readily extend
20 to general certificates. It is not sufficient for a bank to sign a number, with a "\$1" signature, because the bank knows
21 the number it has signed—and can trace the cash. Digital signatures allow the bank to be "blinded" and then sign, and
22 the recipient can remove the blinding factor from the signed number to receive a valid signed number. This ensures
23 user privacy when spending digital cash.

24 2.4 Main Data Flows

25 TO CENTRAL SDI SERVER (not via ISP)

26 Figure 3 illustrates the basic data flows from the client-level SDI proxy and the vendor-level SDI proxy towards the
27 central SDI data warehouse. The vendor-level SDI proxy only knows the pseudonym of a user when it executes
28 sessions with a user and a particular pseudonym, along with what ever basic profile information and other explicit
29 information the user provides the vendor's server with.

30 The client-level proxy server logs all web surfing activity that a user engages in under each pseudonym, but does not
31 have all information about the interaction of each user within a session with a single vendor. The client-level proxy also
32 knows the pseudonyms that are equivalent for a user, and can provide more integrated information than vendors.

33 FROM CENTRAL SDI SERVER (to client and vendor)

34 Figures 4 and 5 illustrate data flow from the SDI data warehouse server to the client-level and vendor-level proxy
35 servers. The updated user profiles are periodically requested by client-level proxy servers (pull), to enable the users to
36 maintain an up-to-date profile for each pseudonym that they operate. The update to the vendor-level SDI proxies occurs
37 according to one of two main modes, as illustrated in Figure 6. A vendor may have sensitive information that it is not
38 willing to release to the central SDI server. In this case, Figure 6 (b) the central SDI server releases profiling
39 information to the vendor-level SDI proxy, where analysis is performed. In Figure 6 (a) the vendor releases all relevant
40 information to the central SDI server, and receives updated profile information.

41 Figures 4 and 5 show the preferred mode of profile releases under our architecture. The Secure Data Interchange server
42 releases pseudonym profiles to users, when verified requests are received from the proxy server that represents the
43 pseudonyms. The SDI server also releases encrypted profiles to the proxy server, either during daily updates, or when a
44 request for update is received from a proxy server. The proxy server maintains an encrypted profile for each
45 pseudonym in its user-base. The profiles are encrypted with a private key of the SDI server. When the proxy server
46 relays a message from a user to a vendor it also augments the message with this encrypted profile, so that the vendor
47 can make use of profile information about the user if the vendor has purchased the analysis ability from SDI. SDI will
48 supply analysis capabilities to vendors, that allow vendors to gain the benefit of analyzing profiles for new users, or

CONFIDENTIAL

- 1 users currently in their user-base, but without violating the privacy rights of users—for vendors can access only the
2 results of the analysis, not the profile itself.

3 2.5 Underlying Cryptographic Infrastructure

- 4 All messages sent between users and vendors can be encrypted to prevent anyone other than the intended recipient from
5 being able to read them. There are many technical solutions to this problem, including asymmetric public key/private
6 key schemes such as the RSA encryption technique (although, users would need a unique key pair for
7 each pseudonym). New "light" security protocols, such as using asymmetric key cryptography for initial validation,
8 followed by shared key encryption are attractive—especially when the asymmetric key infrastructure is inefficient
9 when users need a different key pair for each pseudonym anyway. << cite the new paper from Bell labs here >>
10 We use cryptographic techniques to "digitally" sign messages, in order to validate information contained within a
11 message. A digital signature is computed through encryption with a private key, known only to the certifier, but the
12 signature can be verified with the corresponding public key. This provides a recipient with a high degree of confidence
13 that the message was indeed generated as claimed. An example technology for generating signatures, or "message
14 digests", is MD5.

15 3. The Client-level Proxy

- 16 The client-level proxy, implemented as a client program running on the user's client machine is responsible for
17 managing all data transfer between the client (meaning the client machine and the user) and vendors, or other users. In
18 particular, a key function of the client-level proxy is to implement a pseudonym-management policy for a user—that is
19 able to exert complete control over the ability of vendors to compile data on users. The client-level proxy also
20 negotiates privacy and data-use practices with vendor level proxies. The proxy also allows users to control the addition
21 of demographic and other personal information to electronic profiles, and adds noise to data fields to protect user
22 identity. Figure 2 shows the client-side view of the Secure Data Interchange.

- 23 The client-level proxy maintains profile information for a user's collection of pseudonyms, and allows the user to view
24 and challenge profile information. The proxy also provides a rule-based interface to allow a user to select appropriate
25 privacy/personalization policies.

- 26 The client-level proxy also retrieves pseudonymous e-mail for a user from pseudonymous e-mail boxes, and maintains
27 shared keys for use with each vendor that the user interacts with. Finally, the proxy can personalized generic
28 information provided by a server according to rules provided by a server and profile information stored at the client,
29 and filter/validate incoming e-mail messages.

- 30 The proxy also provides ancillary services, such as automatic user-verification with web pages, a "secure cookie"
31 system to allow servers to maintain stateful interactions with users without allowing the identity of users to be
32 compromised by "flags" that are left on the client and retrieved by another vendor later.

- 33 The primary mechanism that protects the identity of a user across multiple vendors and service providers is the ability
34 to interact pseudonymously with vendors. The user can choose a unique pseudonym for each third party with which
35 she interacts, and be absolutely certain that she is the only party that knows his/her true identity, other than the
36 trusted proxy server where the pseudonyms are registered. There is no way that a vendor can know anything about the
37 transactions that a user has had with other vendors under alternate pseudonyms unless the user chooses to disclose the
38 equivalence of pseudonyms, or use the same pseudonym across multiple vendors.

39 3.1 Initialization

- 40 The Client-level SDI proxy server runs on the user's client machine, and acts as an intermediary between the user and
41 the Internet, intercepting all outgoing and incoming messages. Given that the user runs a standard Internet browser, e.g.
42 Netscape or Internet Explorer, the proxy can be implemented as a plug-in to the browser, integrated directly into the
43 browser, or downloaded as Java (or some-other platform-independent) code. The browser is configured to use the SDI

CONFIDENTIAL

1 proxy as its proxy, and the SDI proxy itself connects through the ISP-level (or other intranet gateway) proxy server to
2 the Internet.

3 A new user must down-load the client-level proxy server that will run on his/her local host machine (or an SDI-enabled
4 browser); and configure his/her browser to connect to the first-level proxy server. Furthermore, the user must be located
5 on an SDI-enabled network, where the Internet Service Provider has a SDI second-level proxy server at the gateway to
6 the Internet. A new user connects to the main Secure Data Interchange through his/her browser (configured through the
7 first-level proxy server), by entering the SDI URL (e.g. <http://www.sdi.com>). The log-in page will prompt for a new
8 user-name and password.

9 3.2 New-user Registration

10 The client proceeds to automatically generate a unique SDI user ID code, and provide information about the user to a
11 central SDI database -- although this information will not be linked to a user's pseudonyms. A flow-chart for the
12 process of registering a new user with SDI is show in Figure 10.

13 When a user first registers with SDI the user provides the client-level proxy with personal information, such as its
14 name, mailing address, and e-mail address (at a minimum). The client-level proxy registers the user with the central
15 SDI server, providing the server with the name, address and e-mail address of the user. Other basic user information
16 could include demographic information, for example a users job, marital status etc. The user can configure his/her
17 client-level SDI proxy to release some of this information automatically to vendors.

18 At this stage the central SDI server must verify the identity of the user, and also check that the user is not already
19 registered with SDI. The method for verifying the identity of a user could include requesting that the user provides
20 his/her social security number, or some other institutional solution that is used for this purpose. In the future we could
21 envisage an electronic system for such an identity procedure, but the method might require for the user to execute this
22 initial step in person with the presentation of a recognized photo ID. The central SDI user ID server maintains a
23 database of all users that are registered with SDI, and checks that that the user is not already registered with the system
24 of secure data interchange.

25 When a new user registers with SDI the user must create a unique public key/private key pair. This key pair can be
26 generated only once for a person, and although the central SDI user ID server does not know the key pair, the server
27 can verify that a key pair is only generated once - because a new user must present proof of identity to establish an
28 account. The unique key pair is used by the client-level SDI proxies to generate new pseudonyms for users, and to
29 verify (when necessary) that a user has only one persistent pseudonym for each vendor.

30 The client-level proxy now generates a unique user identifier, UUID. This is blinded, and signed by the central SDI
31 server so long as the identity of the user can be validated. The technique of blinded signatures is discussed in the
32 Appendix. The client-level proxy now removes the blinding factor, and has a signed UUID that it uses when it is
33 necessary to generate new pseudonyms and request new certificates. The central SDI server that validates new users
34 has a public/private key pair for the purpose of validation, eg. (PKSDI, SKSDI).

35 The central SDI proxy also provides the user with a signed certificate of some universal identifier, such as its Social
36 Security Number, that the user can use to generate other certificates from "User Certifying Authorities".

37 Furthermore, the UUID acts as a public-key, and the client-level proxy also generated a private-key. The client-level
38 proxy can now sign messages with its private-key, and provide the signed to UUID, to verify that (1) the UUID
39 represents a validated user; (2) it is the client-level proxy authorized to act for the user, because it has the private-key
40 associated with the UUID. The technique of public key/private key cryptography is discussed in the Appendix. The
41 client-level SDI proxy uses the private key to authenticate messages that it sends to other modules within SDI, such as
42 Pseudonym administering servers.

43 The unique user ID for a user does not carry any information about the user, its sole purpose is to provide a unique
44 identity.

45 The central SDI server does not know the unique user ID that is associated with a user because we use the method of
46 blinded signatures, but nevertheless a third party or another SDI module receives a guarantee that the user ID has been
47 validated by the central SDI user ID server, and that only one user ID was certified for the user. This is a useful

CONFIDENTIAL

- 1 property because the system of SDI is absolutely and unconditionally secure from security violations that could lead to
2 the security of a users identity to be compromised.

3 3.3 Selecting Privacy and Profile Management Policies

- 4 The next stage in registering a new user with SDI is to establish privacy and profile management policies. We describe
5 a "pseudonym management policy", and a "profile management policy". A user can define how he/she wishes to
6 interact with various classes of vendors (depending on the nature of the business that the vendor is engaged in), the
7 kinds of uses to which the transactional information that a vendor collects can be put to, and the amount of information
8 that a vendor is authorized to release. The user can also specify a "basic profile" that the user is willing to release to any
9 vendor, irrespective of the vendor's policies. The profile-management policy is further broken down into the "data-
10 release" policy and the "use-of-data" policy. The client-level proxy manages a user's interactions with vendors, to keep
11 them within desired policies.

12 3.31 Pseudonym Management Policies

- 13 Pseudonym management, coupled with the ability to add noise to information provided, allows users to exercise
14 complete control over the ability of vendors to collect and exchange information about them. The client-level proxy
15 contains a rule-based interface that allows the user to define a pseudonym-management policy.

16 Abstract Policy Hierarchy:

17 Level 0 (Highest)

- 18 At the highest level of privacy a user chooses to interact anonymously with every vendor, so that vendors cannot even
19 personalize its service to a user over an extended interaction, because the vendor will never know who the user is. An
20 anonymous pseudonym is simply a one-time traditional pseudonym, where a PAS does not need to check that a user
21 does not already have a pseudonym for a vendor.

22 Level 1

- 23 A user can choose to interact with every on-line server under a unique pseudonym. This completely prevents a vendor
24 from knowing anything more about the user than the information that can be inferred from the interaction itself. So
25 long as that information does not identify the user, a vendor cannot access any information from external databases
26 (such as demographic information), and a vendor cannot access any information collected by other vendors about the
27 user. The vendor can still personalize information that is displayed to the user, but only on the basis of its own
28 historical information for that user.

29 Level 2

- 30 ** DP—note that the mode of information release by vendors and users allows users to provide profile information to
31 competitors, but only when the client-side SDI proxy can do a useful job of monitoring details of a transaction (in
32 general this is hard).

- 33 At level 2 a user can choose to share pseudonyms for groups of vendors (for example a user might choose the same
34 identity for all vendors that sell music and books), and receive personalized service from each vendor on the basis of
35 her extended profile.

- 36 The system of SDI allows users to maintain "ownership" of the dossier of information by: (1) providing a different
37 pseudonym to each vendor, but providing the central SDI database with information about which pseudonyms are
38 equivalent; (2) allowing vendors to take advantage of the combined profile when providing information/products to a
39 user by allowing personalization on the client machine but without release of the complete profile for the group of
40 vendors to a vendor.

CONFIDENTIAL

1 ** DP: The client proxy server will release a user's profile (for a pseudonym) to the central SDI server, but
2 anonymously, so that the central SDI server can perform pseudonym level profiling in order to enhance vendor user
3 models, without compromising the privacy of the user. — What about anonymous release of information to the central
4 SDI server? Not sure how this works....

5 Level 3

6 At level 3 the user can interact with every vendor under the same pseudonym, or anonymously with some vendors until
7 enough trust has been established.

8 Level 4

9 At the lowest level of privacy protection (none) the user can simply use a unique ID that is linked to his/her true
10 identity to interact with every vendor. At this level the user has no control over the dossiers of information that can be
11 collected by groups of cooperating vendors.

12 Figure 6 illustrates level 2 - which is the preferred mode of interaction. The client-level SDI proxy for a user submits a
13 request for information to a vendor server, including the pseudonym of the user. The vendor's server can perform some
14 personalization at the server level, based on data that it has accumulated from previous interactions with the user's
15 pseudonym, and also push generic information and rules for processing the information to the client machine, where
16 information is processed according to a user's extended profile for the group of vendors to which the vendor belongs.
17 The rules could be implemented as Java code or Javascript, and the generic information pushed as XML documents, for
18 example.

19 3.32 Implementation Details

20 The central Secure Data Interchange server categorizes vendors that register with SDI, by assigning labels from a set of
21 classifiers, that indicates the business of the vendor (i.e. the services and products that a vendor provides). The set of
22 classifiers might include: music goods, news media, vacation packages, groceries, clothes. We denote the classifiers
23 abstractly as L_1, L_2, \dots, L_N . There is also a label for vendors that are not currently registered with SDI, denoted L_0 .
24 Given the labels, each vendor has an associated set of labels, denoted $L(V_j)$, for example $L(V_4) = \{L_1, L_4, L_{12}\}$.

25 The user is able to configure an appropriate pseudonym-management policy at the first-level proxy level that runs on
26 his/her host machine. The user assigns a pseudonym-management action to different vendor-classes. The set of actions
27 that SDI provides can include, but are not limited to:

28 A. Anonymous interaction with vendors in this class (i.e. use a different pseudonym every time the user enters the
29 Vendors' sites.)

30 B. One pseudonym per vendor in this class, but the same pseudonym for all visits to the same vendor.

31 C. One pseudonym for all vendors in this class, and the same pseudonym at all times.

32 Pseudonym-management action A provides stronger data-privacy than action B, than action C, but action C provides
33 the most opportunity to vendors to exchange information on the user and provide personalized and informed service to
34 the user. This is a basic tradeoff that the user must make - between privacy and personalization.

35 A policy maps a vendor to a "management action", i.e. with Vendor V I will use the same persistent pseudonym for all
36 interactions, but it will be unique to that vendor.

37 The vendor's are grouped according to their "business classifier" labels. For example, group 1 might contain all
38 business that sell books or CDs, group 2 all businesses that sell computer hardware, etc. Each vendor group has an
39 associated management action, e.g. vendors in group 1 (books/CDs) can share the same pseudonym for the user—and
40 exchange information as provided by the user and the vendors to the central SDI server. Other groups could be defined
41 by the labels that vendors DO NOT have, e.g. group 3 could be vendors that do not sell credit cards.

42 Clearly, it is possible for a vendor to be categorized into more than one group, within these rules. For example a vendor
43 that sells books and Compact Disks would be placed in the group of vendors that sells books, and a group of vendors

CONFIDENTIAL

1 that sells CDs. When this occurs the rule-base must choose the most appropriate group, for example on the basis of
2 profiling a group of vendors according to the profiles of their user-base. Alternatively, it is possible to "partition" an
3 individual business into different core businesses that fit cleanly into a single group, and then ensure that a user's
4 interactions respect the current group that the vendor is a representative of.

5 The user can choose a mapping from labels to groups that best fits his/her own privacy needs, and then continue by
6 assigning pseudonym management actions to each group. For example, the user might decide that all vendors that sell
7 flights, books, and music should be assigned the same pseudonym for all transactions (i.e. pseudonym action C), while
8 all vendors that sell financing, loans, and credit cards should be assigned to the same pseudonym for all transactions
9 (i.e. pseudonym action C), but to a different pseudonym than the first class of vendors. Similarly, a user might require a
10 single pseudonym for each vendor that is not classified by SDI (pseudonym action B), and anonymous pseudonyms for
11 vendors that are classified as suppliers of adult material.

12 We allow vendors to provide different service levels, depending on the pseudonym action that a user chooses. For
13 example, a vendor might desire to provide a better level of service to a user that interacts with a persistent pseudonym
14 than a user that interacts with an anonymous pseudonym. The first-level proxy sends a token that certifies the type of
15 pseudonym-action policy that a user has chosen when a user connects to a vendor's site. This is explained in more
16 detail in Section << add section number here >> below.

17 3.4 Profile management policy

18 The basic mechanism that a user has to prevent vendors exchanging information with other vendors about his/her
19 transactions is the pseudonym-management policy. However, there is still a chance that a user's identity can "leak" to
20 another vendor if the user has performed a transaction or provided information that reveals the user's identity. For
21 example, if a user purchases a flight from Philadelphia to San Francisco on UB 004 that arrives in San Francisco at
22 11.50am on 3110199 from one vendor, and then immediately purchases a car rental from San Francisco International,
23 indicating arrival on flight UB 004 from another vendor, then there is a high probability that it is the same user. If the
24 first vendor notifies all major vendors that offer car rentals that it has just completed a sale of a flight ticket to a user
25 with pseudonym C1, and another vendor sells an appropriate car rental to user with pseudonym C121, then the two
26 vendors now have some positive evidence that the two pseudonyms might belong to the same user, and in future can try
27 to cross-sell appropriate products.

28 3.41 VENDOR-SIDE PROFILE MANAGEMENT

29 This is where the role of the "profile management policy" is important. A user can assign a "data-release" action to
30 each class of vendor, depending on the types of actions that a user does with a vendor. There is no way to prevent a
31 vendor releasing the details of a transaction with a user to other vendors, and if the details allow other vendors to
32 identify the user, then the user's privacy can be compromised. However, we can control the type of information that a
33 vendor provides to the central SDI server, for exchange with other vendors.

34 To help with this process, the central SDI server classifies the type of service/goods that a vendor provides according to
35 whether they are "anonymous" high-volume goods such as compact disks or newspaper articles, or "non-anonymous",
36 low-volume, possibly personalized goods, such as cars, flights, or property. A vendor must initially enter a contract
37 with a user about the type of information that it can release to other parties, i.e. the choice of data-release actions can
38 include, but is not limited to:

- 39 A. Release no information.
- 40 B. Release randomized information.
- 41 C. Release all information.

42 The vendor-level SDI proxy is able to randomize transaction information automatically if option (B) is selected, to
43 allow the vendor to submit profile information to the central SDI server without breaking the user's privacy
44 requirements.

CONFIDENTIAL

1 The vendor must provide the central SDI server with its certificate of agreement with the user, when it submits
2 information. All profile information can be verified by the user, and in the randomized section we also describe a
3 technique to validate randomization of profile information. The client-level proxy maintains profile information for
4 each of a user's pseudonyms, and allows a user to view and modify (challenge) profile information that has been
5 provided by vendors and other parties. The proxy requests periodic profile updates from the central SDI server,
6 providing validation via the PID and associated private key that the proxy is authorized to receive profile information.
7 For example, a user might choose that all vendors that sell flights can only release randomized information to the
8 central SDI server or other vendors, while vendors that sell compact disks can release complete information on
9 transactions to the central SDI server and other vendors (although this information will only relate to the actions of the
10 user under the same pseudonym).

11 3.42 CLIENT-SIDE PROFILE MANAGEMENT

12 Orthogonal to pseudonym management is profile management, which determines how a client-level proxy will release
13 profile information to vendors.

14 Level 0 (Highest)

15 Release no information.

16 Level 1

17 Release randomized information. The information is randomized before release to prevent any compromise of the
18 user's identity from profile information that is too specific, while enabling useful personalization at the server level,
19 and enabling useful analysis at the central SDI server.

20 Level 2

21 Release non-randomized profile information for a user under the relevant user pseudonym.

22 We ensure the accuracy and content of profile, given that updates can be made by any third party that is privy to
23 information by allowing users (individuals) to specify privacy constraints that vendors must uphold in reporting
24 information to other vendors and/or SDI

25 3.43 CLICKSTREAM DATA (CLIENT-SIDE POLICY)

26 Clickstream data must be logged at the client, because of proxy caching. This data is released periodically to servers.
27 The client-level proxy server that runs on a user's host machine is in a unique position of being able to monitor the user
28 across different pseudonyms and across different vendors' sites. For example the proxy-server can monitor:

- 29 1. Clickstream data across different vendors' sites and pseudonyms.
- 30 2. Data that is displayed to the user (text, names of graphics objects.)
- 31 3. Input provided by the user at the keyboard of the host machine.

32 We term all of this data "clickstream" data because it is gathered by passively observing the actions of the user, and not
33 by direct question-and-response. The clickstream data policies that are available to a user can include, but are not
34 limited to:

- 35 A. Release no information.
- 36 B. Only release data on the URLs of the most recent sites visited.
- 37 C. Release data about the URLs of the most recent sites visited, and the information displayed to the user.
- 38 D. Release data about the URLs, the information displayed, and the information entered by the user.

39 In addition, the clickstream data policy restricts the depth of information that is provided, i.e. for how many previous
40 sites, and can also restrict the data released to data that was collected under the current pseudonym of the user. The data

CONFIDENTIAL

1 that is released can also be randomized in the same way as explicit data is randomized, for example by removing time-
2 stamp information.

3 3.5 Certificate Management

4 User certifying authorities issue certificates that certify properties about the user that is represented under a particular
5 pseudonym. This service is important to a pseudonymous electronic commerce environment because vendors will
6 require guarantees about certain properties of users that they do business with, such as the age and credit worthiness of
7 a user. Although outside of the basic SDI framework, the existence of such authorities is assumed in the description of
8 the basic operation of SDI.

9 When a user needs certificate C(P) for pseudonym P, a user can use the following steps, to gain a certificate without
10 providing a User Certifying Authority with any new information. The user sends a message to the Certifying authority,
11 containing its SSN certificate, that it received on initial registration with SDI. The User certifying authority can then
12 check the certificate, and if it is valid and the user can be certified for the requested property, signs a blinded message
13 to certify the new property—for a one-time pseudonym specially generated by the user. This certificate can then be
14 unblinded, transferred to certify other pseudonyms that the user holds, and used for certification purposes << See D
15 Chaum's work on digital certificates >>

16 Certifying Authority servers have key pairs, for example (PKCA,k, SKCA,k), which again are using for certification
17 and encryption purposes.

18 The certificate can be trusted so long as the certifying authority keeps its private key secure. To give an example, a
19 trusted third party can exist to certify that a user represented with a particular pseudonym is above 18 years old. The
20 third party can maintain a private key/public key pair (SK_18, PK_18) and sign a message that includes the pseudonym
21 of the user with its private key to generate a certificate, C_18(U), that will assure other parties that the user is 18 years
22 old. This certificate can then be requested by other vendors and information providers, and checked for validity within
23 a public key infrastructure that maintains a faithful copy of the public key of the certifying trusted third party.

24 The first-level proxy server is responsible for certificate management. When a certificate is required by a vendor, then
25 the proxy server checks whether the certificate has been issued to the user under one its pseudonyms. If this is the case,
26 then the proxy server will simply transfer the certificate to the appropriate pseudonym for the user with the current
27 vendor, using a technique taught by D. Chaum << more information here >> There are various certifying authorities
28 within the Secure Data Interchange system, that are able to use a verified social security number to provide certification
29 about the user. When a new user registers with SDI, the central SDI server provides the following social security
30 number certificate to the user, S((SSN, PKSSNC), SKSDI), that links the social security number of the user to a public
31 key that was established for this purpose.

32 When the user needs a new certificate, the first-level proxy server creates a new key pair, using the method described
33 above to certify that this is the first certificate of this kind that the user has requested. The first-level proxy server
34 receives a validated new public key, S(PKP, SKCAS), that it sends to the Certifying authority. The first-level proxy
35 server now generates a blinded certificate number, B(CERTP), and transfers the social security certificate that is signed
36 by the central SDI server to its new pseudonym. The first-level proxy server forms the message M=(B(CERTP), S(
37 (SSN, PKP), SKSDI)), signs it with its secret key SKP, and sends it to the Certifying authority.

38 The certifying authority now verifies the public key, PKP, and then verifies that the user has the appropriate secret key,
39 i.e. checks that the message M is correctly signed. The certifying authority continues by verifying that the user with
40 pseudonym P has social security number SSN. Then the authority determines whether the user with social security
41 number SSN has the correct property, and if it does, signs an association between the blinded certificate number that
42 was provided by the first-level proxy server of the user and the public key of the user's pseudonym.

43 Finally, the first-level proxy server can remove the blinding factor, and it is left with a signed certificate that associate
44 the public key for its new pseudonym, PKP, with the certificate number. The certificate is represented CP = S(
45 (CERTP, PKP), SKCA).

CONFIDENTIAL

- 1 The certifying authority now knows that the user with pseudonym PKP has been issued with a certificate, and
2 furthermore, the authority knows the social security number of that is related to the pseudonym. However, the user will
3 never need to use the pseudonym again, because it can transfer certificates issued in one pseudonym to certificates
4 issued in another valid pseudonym using a method taught by D. Chaum << more details here >>. The Certifying
5 authority therefore gained no information, because it already knew the relation between social security number and
6 property that it is certifying. For example, we can transfer certificate CP to certificate CQ = S((CERTP, PKQ), SKCA).
7 << actually, we need to also change the certificate number, else other vendors can still form a portfolio >>

8 3.51 Pseudonymous Interaction with a Vendor

- 9 The proxy can validate that it represents a particular user under pseudonym P by sending a message to a vendor with
10 the signed PID, and signing a challenge provided by the vendor. The vendor can validate that the signature corresponds
11 to the PID. See the Appendix for a discussion of this challenge-response mechanism.

12 3.6 Generating a New Pseudonym

- 13 The fundamental model of interaction between a user and a vendor under SDI will be pseudonymous. When the
14 pseudonym management policy of user requires that a new pseudonym be generated for a user, it is necessary to
15 validate the new pseudonym to verify to a vendor that the user has a unique pseudonym for its site.

- 16 A user can interact under different pseudonyms with each server, as dictated by her privacy and profiling policies, and
17 the declared policies of vendors with which she interacts. Each pseudonym is associated with a signed PID,
18 (pseudonym ID), a private key that is useful to validate that the PID has not been stolen. The client-proxy also
19 maintains a shared key with each vendor, because symmetric encryption/decryption is cheaper than asymmetric public
20 key/private key encryption.

- 21 There does not exist a central database of public and private keys for users, or user pseudonyms. We use the method of
22 blinded signatures to certify user-generated key pairs, so that only the client-level proxy servers can link PIDs to User
23 identities. Client-level SDI proxies generate new PIDs, that are unique with a high degree of probability << see UUID
24 reference >>, and blinded before authenticated for use within the system of Secure Data Interchange.

- 25 A key feature of the system for administering pseudonyms is that the pseudonym administering authorities cannot build
26 dossiers of the pseudonyms that are authorized for each user, because users submit "blinded" PIDs to be validated. The
27 only information that a PAS has is for each unique user ID, what vendors has the user registered with. This information
28 cannot be linked to the real (physical world) identity of the user because the central SDI user ID server has no
29 information about the user IDs that are authorized for each user that registers with SDI (because the UUID is blinded
30 before signed by the central SDI server).

- 31 To generate a new pseudonym for a particular vendor the proxy requests a new PID from the "Pseudonym
32 Administering Server" that has a trusted-agent relationship with the vendor. The proxy provides the PAS with the
33 signed UUID, and if the PAS can verify that the UUID has not already applied for a PID for use with this vendor then
34 the PAS will sign a blinded PID provided by the proxy, with a signature that indicates the vendor that it is valid for—
35 and make a record that a PID has been authorized for user with UUID and vendor V.

- 36 Each Pseudonym administering server has a public key /private key pair (PKPAS, SKPAS) for each Vendor for which
37 it validates new pseudonyms. A PAS will sign the public key of a pseudonym using the PAS private key associated
38 with the particular vendor. In effect, the operational pseudonym that a user uses is the triple (S(PK, SPAS), SK, IP),
39 representing a pseudonym that is authorized for a particular vendor server.

- 40 There is no information compromise here, other than it becomes possible to construct the set of vendors that a user with
41 UUID has applied for PIDs with. However, the only entity that knows the true identity of the user with UUID is the
42 client-level proxy, because the UUID was blinded before signed by the central SDI server when a new user is
43 registered.

- 44 In this way it is possible to provide vendors with guarantees that PIDs are once-in-a-lifetime for a user, so that vendors
45 can continue to achieve at least the same level of personalization as they do without SDI in an SDI system. In this way
46 SDI can guarantee once-in-a-lifetime pseudonyms to those vendors that require persistent interactions with users.

CONFIDENTIAL

1 The PID and associated private key can be used to validate an initial information exchange with a vendor, but it is also
2 possible to perform follow up message exchange using a shared key pair—this is more efficient to implement than an
3 asynchronous key pair cryptographic solution. Messages can be encrypted with the shared key, that only the user and
4 the vendor know. This (1) validates that the message is from the sender; (2) ensures that only the intended recipient can
5 read the message. << refer to new Bell Labs paper here >>

6 3.61 Implementation Details

7 Figure 11 shows a flow-chart for creating a new "relationship" with a vendor. The Pseudonym Administering servers
8 (PAS) provide for validation of new pseudonyms. Each vendor selects a PAS that will be responsible for managing
9 pseudonyms for its domain. When a client-level SDI proxy requires a new pseudonym for a user UUID, the proxy does
10 the following steps:

- 11 1. Request the URL of the PAS from the server for the new vendor that the user wants to initiate a persistent
12 relationship with.
- 13 2. Send a message to the PAS with the tag "New Pseudonym", the URL of the vendor, and its validated UUID. Also
14 be prepared to answer a challenge/response with the private key associated with UUID.

15 The PAS for the vendor then checks that it has authorization to administer pseudonyms for the vendor. The PAS
16 performs the following steps:

- 17 1. Check that the vendor URL corresponds to a vendor for which it has pseudonym-management authority.
- 18 2. Verify that the W I D is correctly validated for use within SDI (indicating that the user is a member of SDI).
- 19 3. Verify that the client proxy server represents the user, through a challenge-response sequence that requires that
20 the proxy has the private key for UUID.
- 21 4. Look up the user ID in the database of user IDs that the PAS maintains to check that a pseudonym has not
22 been authorized for this user and this vendor.
- 23 5. Send a message to the user, indicating either (OK, or DENY).

24 Given that the client-level SDI proxy receives an OK response, the client-level SDI proxy then generates a new PID,
25 and private key, and blinds the PID for validation:

- 26 3. Generate a new key pair, (PID, SK).
- 27 4. Blind the public key, and send a message to the PAS with the tag "Request Certification", the URL of the vendor,
28 the signed unique user ID, and the blinded public key, B(PID)
- 29 5. Receives a signed copy of the blinded public key in return, and removes the blinding factor to obtain S(PID,
30 SKPAS).
- 31 6. Generate a new pseudonym, from the components, (S(PID, SKPAS), SK, EMAIL), where EMAIL is a new e-mail
32 address for the pseudonym.

33 The public key that represents the new pseudonym is signed with the private key of the PAS that relates to the
34 particular vendor. This enables a vendor that receives the pseudonym to validate that the pseudonym is unique for the
35 user, to enable persistent interactions across multiple sessions.

36 3.7 Personalization of information/Manage e-mail

37 The client-level proxy can provide vendors with "certificates" to enable them to send e-mail to the user (under a
38 particular pseudonym). Outgoing e-mail includes a certificate that a vendor can use to reply, and a pseudonym ID in
39 place of the standard "from" field. The certificate is of the form S(M, SKP), where P is the pseudonym of the user, and
40 M = (PKP, PK*V). The vendor includes the signed certificate when replying to the user.

CONFIDENTIAL

3.8 Automatic release of personal information

- 2 Maintains certificates as well.

3 3.9 Enhancing profile information

- 4 Receives update requests from SDI server, allowing users to access any profile information that is stored in the SDI
5 server, and request changes to that information if it is not accurate, or if the user does not want content to be available
6 to third parties. This information is only released to users via the proxy server that is authorized to represent the user for
7 the relevant pseudonyms.

8 3.10 Manage iAmWorthIt module/Negotiation

- 9 The proxy server determines what kind of site is being visited, from SDI credentials that are provided to a vendor's
10 SDI proxy server, and either embedded in web documents, or provided in prenegotiation. The proxy then determines
11 an appropriate type of interaction level with the vendor, depending on the profile and credentials of the vendor, the type
12 of information required by the vendor, and the user's privacy policy.

- 13 The system of Secure Data Interchange allows user's to receive compensation, what we will call "community dollars",
14 in return for providing information to vendors. The architecture supports a negotiation between user's proxy servers
15 and vendor's proxy servers, to strike a deal about information use and compensation. In some cases the exact nature of
16 an offer may not be anticipated, and the user can be contacted directly.

- 17 The vendor's host level proxy server is then granted the permission by the client level proxy to receive certain
18 appropriate information. The vendor can use profile information about a user (for a particular pseudonym) to present
19 appropriate services, products, and prices, including custom priced items and promotional offers (as suggested in the
20 co-pending application entitled "System for Customized Prices and Promotions").

21 3.11 Remote Retrieval of Profiles

- 22 The client-side SDI proxy is designed to be configurable from a remote database. This enables a user to maintain a
23 persistent SDI profile across different client machines, for example at work and at home. The profile, pseudonym and
24 key information that represents a user in its interaction with SDI-enabled vendors and information providers during a
25 session can be saved, and then encrypted and stored in a remote database for user name and password access.
26 Alternative technologies include smart card techniques, where the data is stored on a portable device in encrypted form.

27 3.12 Use-of-data Policy (Access Control Policies)

- 28 The Secure Data Interchange system also allows a user to place constraints on how the vendor can interact with the
29 user, and also on what uses can be made of the information that the vendor releases to the central SDI server and other
30 vendors in the system. The user can specify whether or not the vendor can send electronic and/or physical solicitations
31 to the user, and the user can also specify whether data that is released can be used for (any of): solicitation by other
32 vendors, personalization of service should the user visit a vendor's site. For example, a user might require that any
33 information that is released by a vendor to the central SDI server, and then possibly exchanged with other vendors, is
34 only used to personalize the service and products that a vendor offers to the user should the user visit the vendors site,
35 and is not used for electronic or physical solicitations. The use-of-data policy is augmented to the pseudonym-
36 management policy, and defined over classes of vendors.

- 37 SDI enables a user to allow limited and secure data sharing between vendor types of its choice. i.e. if I use the same
38 pseudonym for multiple vendors, I am still not saying-OK you can all share my data (although I might choose to say
39 that). I can say-OK, you can use my data to personalize the service that you offer me without knowing what I did with
40 other vendors. Also, vendors themselves can say-here is the data, allow other vendors who are not competitors to
41 access the data, and refine their models, but not actually access details.

CONFIDENTIAL

1 << page 110 -- 113 1995 patent the user can issue strict guidelines to SDI, vendors, and his/her proxy server, about
2 how information that is stored can be used. >>

3 3.13 Dynamic Privacy Management

4 When a user clicks to a new URL, then the first-level proxy server first checks its local cache, to determine whether it
5 already has the vendor's classification certificates. The proxy server requests the classification certificates and public
6 key certificate from the vendor if the vendor is not already in its cache. The proxy server continues by verifying the
7 integrity of the certificates, and checking that the vendor matches the ID enclosed in the public key certificate.

8 The proxy server continues by looking up the pseudonym-management policy for the vendor, and checking whether an
9 appropriate pseudonym already exists in its local database. If the vendor belongs to a new class of vendors that do not
10 yet have a pseudonym, or the vendor belongs to a class of vendors that require one pseudonym each, or the vendor
11 belongs to a class of vendors that require anonymous interaction, then the first-level proxy server continues by
12 generating a new pseudonym.

13 When a user interacts with a vendor under a persistent pseudonym then the user must have the pseudonym certified by
14 a Pseudonym Administering server, to certify that this is the only pseudonym that the user has registered for the
15 vendor. Each vendor selects a Pseudonym Administering server that will certify all of its pseudonyms. If the user
16 already has a pseudonym for the class of vendors that the vendor belongs to, but the pseudonym has not been certified
17 for this vendor, then the first-level proxy server will have the pseudonym certified. Similarly, if the user requires a new
18 pseudonym for this vendor, then the first-level proxy server will generate a new key pair and have the public key
19 certified by the Pseudonym Administering server.

20 The first-level proxy server performs the following steps (we use the technique of blinded signatures so that the
21 Pseudonym Administering server does not know the public key that it certifies):

- 22 1. Generate a new key pair, (PKP, SKP) - or - lookup an existing
23 key pair that will form a pseudonym for this vendor.
- 24 2. Blind the public key, by multiplication with a random number,
25 and form the message $M = (S(PK * U), SKSDI), S(B(PKP), SK * U), S(PK * V, SKSDI))$.
- 26 3. Encrypt the message with the public key of the Pseudonym
27 Administering server, and send the message.

28 The Pseudonym administering server checks that the unique public key of the user is signed correctly by the central
29 SDI server, and that the user has the private key associated with that public key by using the public key to verify that
30 the blinded pseudonym public key is correctly signed. Finally, the Pseudonym Administering server checks that the
31 public key of the vendor is signed correctly by the central SDI server, and then checks that the user does not already
32 have a pseudonym for the vendor, in a database that it maintains of which unique user PKs have requested pseudonyms
33 for each vendor.

34 If everything is OK, then the Pseudonym administering server signs the blinded public key for the user's new
35 pseudonym using a key pair that it maintains for the vendor, and returns the signed blinded key, $S(B(PKP), SKPAS, V)$
36 to the first-level proxy server.

37 The final step in the protocol is for the first-level proxy to remove the blinding factor, so that the user now has a new
38 key pair, (PKP, SKP), with the public key certified by the appropriate Pseudonym administering server private key,
39 $S(PKP, SKPAS, V)$, to demonstrate that this is the only pseudonym that the user has for the vendor.

40 The first level proxy server now connects to the vendor in secure-SDI mode, by sending the vendor the certified public
41 key for the user under this pseudonym that will identify the user to the vendor. The proxy server continues by sending
42 the data privacy policy for the user with this vendor, signed with the private key for the pseudonym. This serves three
43 main purposes: (1) it demonstrates to the vendor that the user is indeed the owner of the public key PKP, because the
44 user has the secret key that is associated with it; (2) it informs the vendor about the data-privacy policies that the user
45 requires; (3) whenever the vendor submits information about transactions with this user to the central SDI server it
46 must also submit this certificate to verify that it is following the user's guidelines.

CONFIDENTIAL

3.14 Dynamic Profile Management

- When a user connects to a site and provides a certified public key, the first-level proxy server also provides a time-stamped certificate of connection, $S(PK*V, T)$, SKP), where T is the current time, PK*V is the unique public key of the vendor, and SKP is the secret key of the user for the pseudonym that it uses with the vendor. This "connection certificate" is used by the vendor to request a profile-release from the central SDI server.
- In addition, the first-level proxy server also provides a basic-profile to the vendor. The basic profile for a user contains no identifying information, but can contain whatever general information the user is happy to release across all pseudonyms, such as the user's age, nationality, state, sex. This is the basic profile that is configured by the user during initial registration with SDI.
- The vendor provides this certificate to the central SDI server, and requests profile information about the user with pseudonym P that is stored within the central SDI server, and is authorized by the user to be released. If the vendor is authorized to receive dynamic profiling information, such as the recent web footprints of the user, and the material that the user has been reading, and the physical location of the user, then this information is released to a vendor when the user connects to a new site, according to the dynamic profile policy of the user.

3.15 Merging off-line and on-line data

- Figure 12 shows a flow-chart for how the central SDI server can request that new information be merged with a user's on-line profiles.
- The central SDI server can associate off-line information about a user with a user's on line pseudonymous profile, even though the central server does not know the user's pseudonym IDs. This can only be done with the user's consent, and may also involve appropriate compensation. Within the system of iamworthit (see section XX) we can credit users for both off-line and on-line information.
- Merging a marketing database with SDI user-profiles can be useful both to initialize the database, for example when asking a user questions to generate an accurate user profile rapidly and efficiently. Off-line data can also add useful richness to on-line profiling information, which may be largely contextual and low on details/factual information. For example, off-line data can include information such as whether a user owns a car, rents an apartment, has house insurance, life insurance etc.
- SDI can also extrapolate correlations to other user profiles, on the basis of common SDI-profiles, for example using statistical techniques.
- It is often the case that individual customers appear in some databases, but not in others. Under normal circumstances, an analyst working across different databases would be faced with a large number of incomplete customer records, each with gaps corresponding to the fields of the databases to which they don't belong. A solution to this problem is offered by SDI, which is capable of drawing correlations between different databases - this information can be used to generate predictions to fill in the gaps of incomplete customer records. The result is a full set of customer records that can be meaningfully sorted or filtered by any of the combined fields, and which can now be handled as a unified set of data, suitable for use by standard database analysis systems.
- In a typical example, SDI might be used to combine a demographic database, such as the one offered by the Econometrics Corporation, with a commercial database, such as the one offered by Claritas. The Econometrics database consists of 180 million different customer records, but at a fairly coarse-grained level of detail, consisting of such information as age, gender, family status, location (at the state, city, or zip code level), and personal income. In comparison, Claritas offers a smaller base of customers, but includes information of arguably higher quality, since it breaks customers down to the geocode (sub-neighborhood) level, and includes much more detailed information on personal spending habits across hundreds of different purchase categories. A logical reason to combine these databases would be to supplement information about customers in the vastly broader demographics dataset with particular predictions about their personal preferences and likely commercial spending habits. One could imagine using this augmented data set to support a web site that instantly customizes itself to new visitors' preferences. Since the number of records in the Econometrics database is equivalent to roughly 72% of the population of the United States, it is likely

CONFIDENTIAL

1 that most first-time visitors to the site will already have a "thumbnail sketch" in the system, and can thus be greeted
2 with an page appropriately configured to their personal tastes.

3 The technical details of the combination process (which have been described elsewhere in the patent) to a large degree
4 depend on the amount of overlap between the databases, that is, the number of customer records which are shared in
5 common.

6 Suppose the demographic databases' fields are coded (x_1, \dots, x_n) , and the commercial databases' fields are coded (y_1, \dots, y_n) . Suppose further that customers in set A appear only in the demographic database, customers in set B appear only
7 in the commercial database, and customers in set C appear in both.

8 The process of supplementing the fields of customers A depends completely on the derivation of the distribution $f(y_1, \dots, y_n | x_1, \dots, x_n)$, which describes the correlation of fields in the commercial database on fields in the demographic
9 database. As previously discussed in the patent, different techniques may be used to create this distribution, depending
10 on the size and variety of C.

11 As a concrete example, one could imagine that set C includes customers from rural areas. The demographic database
12 would reveal that, although their incomes aren't huge relative to the national average, they tend to spend a lot of it (i.e.
13 are active consumers), have large families, and purchase large vehicles. The commercial database might show that they
14 enjoying hunting magazines and Ford trucks. If they live inland, they buy hunting equipment, if they live near the
15 ocean, fishing equipment.

16 If these trends are dominant in set C, they will impact the distribution function. Thus, when a browser from a small
17 town in Texas with a typical income pattern visits the automated website, he could be greeted with discounts on truck
18 accessories and a small sidebar with news on the hunting season. On the other hand, a visitor from a small town in
19 Maine might be given the same truck discounts, but would have news on the fishing season.

20 Although the demographic dataset is arguably the weaker of the two in terms of content, the fact that it contains even a
21 small amount of information on most people in America makes it very valuable for handling first-time visitors, since
22 most of them will appear in it. By using SDI to leverage the more detailed information in the commercial database, we
23 are able to supplement the rough demographic data with predicted commercial preferences. This allows us to construct
24 more detailed thumbnail sketches for each customer, allowing our reception of first-time visitors to be much more
25 appropriate (since knowing personal hobbies or interests tells us much more about a person than general income level).

26 4. The ISP-level Proxy Server

27 The ISP-level proxy server is positioned just behind the firewall of the user's local dial-up network (ISP or Intranet).
28 The proxy provides protection for users operating under pseudonyms from point-to-point attacks and HTTP header-
29 tracking by stripping HTTP header-information and forwarding HTTP packets on to their destination with no
30 information other than their source at the ISP-level proxy server. The ISP-level proxy also supports pseudonymous e-
31 mail, between users, and between users and vendors.

32 Figure 2 shows a couple of users connected to clients, that are in turn connected to the Internet through a local intranet,
33 such as the network of an Internet Service Provider (ISP).

34 The proxy "washes" outgoing messages of any information that would compromise a user's pseudonymity, for example
35 the "referral" field that contains the previous URL of a user in a HTTP message. HTTP messages also leak other
36 information, for example browser software on a user's client machine, the operating system and a user's IP address.

37 4.1 Support for Pseudonymous electronic mail

38 A user can receive electronic mail through the PID and associated IP address of the ISP-level proxy server.

39 The preferred implementation of this system allows the user to periodically check for new mail. The client-level proxy
40 gains access to the mail box that is associated with a pseudonym by providing a correct response (signature) to an ISP-
41 generated challenge. Notice that with this solution, the ISP-level proxy has no way to connect the pseudonyms of a

CONFIDENTIAL

1 user, so long as the user's client is not identified in its messages to the ISP-level proxy server other than by the PID that
2 the proxy makes a request for.

3 ** DP. Must be careful to "wash" HTTP at the client-level proxy as well % as the ISP-level proxy.

4 We can extend this mechanism using a technique taught in the Lucent Personalized Web Assistant. The LPWA
5 provides for a sequential access mechanism to the mailboxes that belong to a user through a one-way function that
6 takes the user's SDI log-in name and password, and an integer from 1 to N, and computes the mailbox location. The
7 mail server does not need to maintain a list of pseudonyms for each user, because the user is able to efficiently access
8 all of its mailboxes sequentially as a function of other information.

9 Another variation, that relies on the user placing trust in the ISP-level proxy server, provides the ISP-level proxy with
10 the e-mail address for each pseudonym. This push method is more efficient, because the ISP proxy and the client proxy
11 communicate only when new messages arrive, but provides the ISP proxy with information to compute all the
12 pseudonyms for a single user — probably undesirable.

13 4.2 Support for anonymous profile-based search

14 The anonymous profile-based search allows a user to release her profile as an addition to a query term to an general
15 search engine (such as Altavista or Yahoo), and have SDI perform additional filtering of the results of the search to
16 refine the pages returned on the basis of their profiles and the user's profile. This is an example of how SDI leverages
17 existing Internet technologies. Figure 9 shows how the anonymous profile based search can be implemented on the
18 ISP-level SDI proxy.

19 4.3 Maintain User Profiles

20 4.4 Support Electronic Payment/Physical mailing solutions

21 5. Vendor-level SDI Proxy

22 The vendors are represented with servers that provide an on-line electronic commerce store-front for the services or
23 products that a vendor provides to users. The servers may connect through their own pseudonymizing proxy servers,
24 but more likely a vendor will not require pseudonymous interactions with users, and will be happy for users to build
25 dossiers on interactions with the same vendor. The vendors make use of the support for pseudonymous electronic
26 commerce functionality that the Secure Data Interchange system supports - for example, pseudonymous payment
27 mechanisms and physical mailing of products and letters to pseudonymous users. The vendors subscribe to the services
28 of SDI, and can receive payments for user information that they provide to the interchange.

29 Vendors connect to users and other vendors through SDI proxy servers. The vendor-level SDI proxy can profile users,
30 periodically release information to the central SDI database, and support various statistical generation modules for
31 analysis of profiles collected and requested from other sources. Figure 1 shows a typical vendor that is connected to
32 SDI through a server and vendor-level proxy server.

33 5.1 Profile Building

34 The main function of the vendor-level SDI proxy is to build a deep profile of users as they interact with a vendor

35 5.2 Profile Release

36 and manage the release of information to the central SDI data warehouse. Figures 3, 4, 5 and 6 show patterns of data
37 flow to and from the vendor.

CONFIDENTIAL

5.3 Request Profile Updates/ Cache metatags/ Manage target object profiles.

5.4 Markup Web documents with Profile information and Personalization commands

A new vendor may choose to operate under different pseudonyms for different users or classes of users, but in general we will assume that a vendor has a unique key pair, and that it is happy to assume the public key of its key pair as an identity, in addition to its more traditional corporate identity. In what follows we assume that the vendor will operate under a single public key with all users, although the extension to multiple pseudonyms is a trivial extension, just using the same technology as described for users.

When a new vendor wishes to register with SDI, the vendor configures a first-level proxy server to manage its SDI functionality. The first-level proxy server is similar to the proxy server configured by new users. Initially the proxy server will generate a new key pair for the vendor, (PKV, SKV), and submit the public key to the central SDI server for signing, along with an identifying code, such as the IP address from which the vendor operates, ID. The vendor receives $S(PKV, ID)$, SKSDI in return. In addition, the vendor is audited by SDI and is certified according to the class of products or services that it provides, receiving message $S(M, SKSDI)$, where $M=(L(V), PKV)$. The vendor is also certified according to the anonymity of the service or goods that it provides, receiving either an anonymous certificate (e.g. if it deals in high-volume, non-personalized goods), $S(M, SKSDI)$ where $M=(PKV, ANON)$, or receiving a non-anonymous certificate (e.g. if it deals in low-volume, personalized goods), $S(M, SKSDI)$ where $M=(PKV, NON_ANON)$. The vendor will provide these three certificates whenever a user connects to the page and the user does not already have the certificates cached for the vendor.

A vendor's data privacy policy can also usefully include the placement of restrictions upon any portions or all of their data disclosed to the interchange, such as which other vendors (or vendor types) are entitled to receive the advertising benefits and what types of advertising benefits derived from his/her data are they entitled to; for what advertising purpose can an "entitled" vendor use it for (if it is not deemed "confidential" the disclosing vendor may request beforehand a copy of the solicitation to his/her customers); limiting use of its data exclusively to other vendors who through their data contributions can provide significant reciprocal benefit to them (often occurring in dealing with competitors); physically confining the data to the interchange (without transferring it to the other vendor); data may only be transferred to non-competitors; advertising solicitations that result from the data is of a form that is non-competitive; recipients of data must remain non-competitors.

For example, a vendor can request that the commercial domain of the requesting vendor not match certain vendor types. Vendors must register their business type with the data interchange, and violators can be punished. Alternatively, data can be constrained only for vendors that are improving existing data models, but not to vendors that are seeking new customer prospects. Disclosing vendors may also require that their advertising campaign and possibly even their corporate identities remain secret. Because SDI is a comprehensive repository of statistical data (and customer rating information across a variety of criteria) about sites, it is able to deploy data mining techniques, provide highly detailed statistical information and verification as to the nature of each vendor's business and its customers. Thus helping SDI to better identify, reclassify and enforce this data privacy policies. Provided that the vendor has agreed to release such information to another type specific vendor, this information can serve as additional useful information (particularly quality and other rating credentials about vendors) to help vendors do a better job of selecting vendors by which they establish their data privacy policies.

6. Central SDI Server

Also shown in Figure 1 is the central SDI server, "SDI data warehouse server". The central SDI server maintains extended profile information for each user, at a pseudonym level, and profile information for each vendor. The central SDI server also supports cross-vendor and single-vendor personalization tools, such as multi-attribute collaborative filtering techniques. The data warehouse operates on a trusted and secure server. The server is able to provide information to vendors that enables models to be enhanced, and synergies between data sets realized, often without providing vendors with direct access to the actual user profiles (unless the user and submitting vendor has authorized such a disclosure.)

CONFIDENTIAL

- 1 The central SDI-server maintains profile information for each pseudonym of every user registered with SDI, aggregated
2 from information provided by multiple sources (e.g. vendors and clients). The server contains modules to perform data
3 mining and collaborative filtering on profiles, and release results.
- 4 The SDI server ensures the integrity of data, and prevents data being used for unauthorized purposes. The server
5 performs collaborative filtering analysis on the large data sets, and releases vendor-specific model enhancements. The
6 central SDI server will also provide targeted mailing to a virtual mailing list of targeted users, without releasing the
7 pseudonyms to a vendor. The main SDI server also maintains an encrypted portfolio of pseudonyms and personal
8 information for each registered user, and is configured to release that information to client-level proxy servers when a
9 user logs-on to SDI. Another important roles of the SDI server is the categorization of vendors' business interests and
10 data-handling policies.

6.1 Profile Updates

- 12 Figure 3 shows modes of data flow to the central SDI server.
- 13 SDI will accept updates to profile for pseudonym P by communication from the client-level SDI proxy server
14 associated with pseudonym P. A user can log new information about one of its pseudonyms by sending a request R,
15 with profile updates, to the central Secure Data Interchange server. The user's identity is verified using public-key
16 cryptography, managed at the client-level proxy. A user can also reveal the equivalence of a set of pseudonyms to SDI.
- 17 The Secure Data Interchange server will accept updates to the profile of a pseudonym from any party in the system that
18 has registered with SDI. The user can request to see the profile for any of its pseudonyms at any time, to check that the
19 information accurately reflects his/her behavior, interactions and transactions with vendors and other parties. A user
20 request for a profile is initiated by a message M from the user. The message M contains a one-time key D, with which
21 SDI will encrypt the profile. Only the user knows the key. The message M is encrypted with the public key of SDI, and
22 then sent to SDI via the user's trusted proxy server.
- 23 SDI verifies that the request for profile disclosure has been received from the authorized proxy server for pseudonym P,
24 and then sends the profile, encrypted with the one-time key D, to the user via the proxy server.
- 25 SDI records a vendor ID when a party other than the user makes an update to the profile of a user's pseudonym. The
26 vendor must be registered with SDI, and make a signed commitment that it will only register information that is
27 accurate and collected in good faith. Whenever a user challenges information, SDI can trace the vendor that provided
28 the information.

6.2 Maintaining Data-use Guidelines

6.3 Use of Data

- 31 Vendors can check in data to SDI in any of the following forms:
- 32 i. Data or Randomized Data, with/without pseudonym
33 ii. Aggregate information on profiles in user-base, with/without pseudonyms
- 34 Suppose that a user has requested that a vendor does not reveal any information about the transactions that a vendor has
35 executed with the user (under a pseudonym). The vendor can still reveal randomized information that will in no way
36 identify a user (even though the user interacts under a pseudonym). We can perform system-wide collaborative filtering
37 on randomized information (details below).
- 38 SDI can reveal information to vendors in a number of forms:
- 39 i. Data or randomized data, with /without pseudonyms
40 ii. Analysis results: For example
- 41 if a customer is interested in product X, suggest product Y
42 if a customer with profile type A hits your site, suggest

CONFIDENTIAL

product Y

if a customer with profile type A is interested in product X, suggest product Y

if customer P hits your site, suggest product Y

if customer P hits your site, and is interested in product X, suggest product Y

The SDI server can provide analysis based on the profile of a user with pseudonym P without requiring that the vendor can access pseudonym P by providing a secure function evaluation procedure, that analyses the encrypted profile and provides results (details below).

The SDI server does not actually reveal any private information about the user to the vendor, so long as the results of the analysis are restricted to recommending products that the vendor has that might be just what the user is looking for. We protect the privacy of users while providing personalization. The options above provide for a whole range of personalization, from very coarse-grained, to very fine-grained, but it is all dynamic and secure. Furthermore, we can provide the ability of a vendor to personalize service to users that hit his/her site, but prevent the vendor from soliciting users.

SDI can also retain control over information on users by placing cryptographically enforced "expiration dates" on the ability of vendors to analyze profiles that are delivered with requests from users. This allows sophisticated pricing models (see later section on pricing models).

There is a hierarchy of personalization that is available to each vendor in the system. As a base level, a vendor can request that SDI performs data analysis on its current user base, and with only transaction information collected at its own site. The next level performs analysis on the vendor's user base, but also drawing on profile information from across multiple vendors' sites. The next level allows vendors to personalize service to incoming users on the basis of an analysis of the profiles of the new users. Finally, SDI may provide a list of pseudonyms and recommendations, so that the vendor can target new users.

6.4 Aggregating Profile Information from Clients and Vendors

6.5 Releasing Anonymous Profile Information for SDI-enhanced Search

Figure 9 illustrates how anonymous profile-based search can be performed on ISP-level SDI proxies, with information provided from the central SDI database.

6.6 Data Mining/Collaborative Filtering Functionality

One of the core purposes of SDI is to provide a common location and format for information that has been gathered from a wide variety of sources and that might require different sorts of analysis. Since its framework is designed to generically handle different types of data and algorithms, SDI can be used as a platform to explore and exploit the rich connections that potentially exist within and across the databases of different vendors and customers.

This section describes in greater detail the types of data, analytical methods, and forms of validation that are available to SDI. The Secure Data Interchange architecture in its preferred implementation integrates the architecture issued U.S. Patent No. 5,754,939 "System for Customized Identification of Desirable Objects" into a system for secure data exchange between multiple parties. The aforementioned patent teaches a method for profiling objects and users over a bi-directional distributed network, such as: an ISP, multiple ISP networks, a Web hosting network, or server software (such as data mining or recommender software) that is linked to a coalition of sites (such as a portal or Internet mall). The current invention, the system of Secure Data Interchange, allows correlations to be identified between vendor's data sets, that allows accurate profiling through the application of statistical methods, without providing vendors with explicit access to the profiles of users — because profiles are provided in anonymized and randomized forms. There are less efficient methods that can be used to identify correlations, for example using customer demographics, and vendor categories, to suggest which vendors might be well placed to form dynamic syndication relationships. With SDI it is possible to leverage as many data sources as are available, about users and the

CONFIDENTIAL

1 target objects with which they interact. In fact the degree of the measure of improvement in predicting user behavior (or
2 increasing clickthrough) is approximately in direct proportion to the square root of the number of user profiles and
3 target profile interest summaries which are known. The emphasis in the aforementioned patent is on the bilateral
4 relationships between vendors and users, and the architecture is not designed to support secure and privacy-protected
5 data interchange and analysis across the user bases of different vendors.

6 In the system for SDI we push control of the profile for each user to the client software that runs on the machine local
7 to the user, and provide for personalization through dynamic processing of information on the user's client machine.
8 We enable vendors to exchange data sets only to the degree that is mandated by users, and provide technical solutions
9 to enable significant leverage of data while maintaining user privacy.

10 With SDI we enable second-level proxy servers for vendors to interact, and also make available the user profile
11 information that is collected at the client-level and ISP-level SDI proxies to develop more accurate profile information,
12 through the combination of "deep" vendor-specific information with broader network-level information. The ISP-level
13 SDI proxies can also capture some information (in a privacy-protected form) about users and vendor sites that do not
14 subscribe to the SDI service. The goal of the architecture is to continue to allow relevant filtering to be applied at sites
15 that a user has never even directly visited before, or interacted with before, without that site accessing the profile
16 information of the user.

17 The supporting architecture as stated in the above referenced patent also allows for profiling statistics to be collected
18 and processed in such a manner that the information and vendor servers may both contain and implement the modules
19 for profiling in a distributed manner. In the present invention the profile generation capabilities are implemented at
20 various levels - in particular, at the second-level proxy servers of ISP networks and vendor servers, and within the
21 central SDI data warehouse.

22 Web page tags include profiles of target objects, user quality ratings based upon overall quality as well as other criteria
23 (e.g., value, price, entertaining, informative graphic/visual appeal, etc.) and data mining analysis of each criteria as it
24 corresponds to target objects, target object attributes, user attributes (and vice versa), trend analysis statistics, location
25 data (for target objects representing physical or geographical items). User information, in addition to profiles, can
26 include data mining and trend analysis statistics, user provided ratings for target objects, and resolution credentials.

27 Vendors can use queries, data mining, menuing, and other techniques (such as described in the aforementioned patent),
28 in order to gain access to desired user profiles that the vendor is allowed access to. The information that is stored
29 against a user's pseudonym in the central SDI server is checked against the data-release certificate that the user
30 provided to releasing vendors, and is therefore data that the user has indicated can be stored on his/her behalf.

31 6.61 Structure of the Central SDI database

32 The central SDI server is structured as a relational database, with data that is submitted by vendors and users to the
33 server indexed by pseudonym. The basic structure of a data record relates a pseudonym ID to a vector of numbers,
34 representing the profile for that pseudonym. (Pseudonym-id, Profile). We also add tags to indicate the usage-
35 restrictions on the profile information, and whether or not the data is randomized.

36 There are several types of information which can characterize both users and items. SDI is intended to function as the
37 intermediary between a vast web of vendors, on the one hand, and individual consumers, on the other. The major
38 sources of data used by this system are therefore:

39 1) Demographic. Such data will most likely be elicited by SDI from vendors and consumers when they initially register
40 for the service, and details very general characteristics about them. It will consist of numbers and categorical values
41 (age, zip code, sex, level of education, etc.).

42 2) Commercial. This is the kind of data any that vendor collects in the course of doing business (especially e-
43 commerce); generally, it links customer codes to purchase items, dates, quantities, and prices. Depending on the nature
44 of the business, this data could be fairly complex, and might well include text. For example, one could imagine that a
45 bookstore, in addition to keeping track of its sales history, collects book reviews, author profiles, and plot summaries.

CONFIDENTIAL

3) Behavioral (vis-a-vis the Internet). When a customer accesses the Internet through an Internet Service Provider (ISP), his ISP is in a unique position to monitor and record every behavior he exhibits while browsing the World Wide Web, writing email, and uploading/downloading files. This information is fairly complicated, as every site visit involves the creation or consumption of content (with associated sound, images, and text), the generation of queries (when using text-based search engines), and the navigation of hyperlinks. An additional complication is that such data is dynamic, in that differing amounts of time are spent for on-line activities and might well reflect the customer's personal preferences (e.g., a long time spent lingering between two choices in an on-line catalog might indicate an equal level of preference, whereas a speedy exit from a site might indicate general lack of interest).

Further information is provided by XML tags attached to pages browsed by the customer. The mere presence of such tags allows for correlations to be drawn between different web pages (e.g., a common XML tag used by travel-related sites), because it implies similarity. Furthermore, it is conceivable that such tags could encode more refined measures of a web page's content, such as browsers' evaluations of its value. For example, a web page of interest to scale modelers, in addition to having images and text related to model trains, might have an XML tag that shows that other scale modelers have given the web site a "five-star" rating. This page should therefore be given a greater weight when SDI is used to create correlations of interest to model hobbyists.

To fully represent these different sets of information, SDI handles collections of data of the following types:

- 1) numerical (e.g. an age, price, or period of time)
- 2) categorical (e.g. a color or musical genre)
- 3) text

A common task for SDI is to compare and correlate different customers, which might well be represented by mixed collections of numbers, categories, and blocks of text. This is handled by treating each customer c_i as a vector in a space whose coordinates correspond to the fields of data available. In the following description we refer to a customer, but when a user interacts with a vendor under a pseudonym, the profile information will only relate to information provided to the central SDI server for that pseudonym.

If there are m numerical pieces of data available, there will be n corresponding coordinates in the data space, (x_1, \dots, x_m) .

For each category i , there will be a corresponding number of values, n_i . Hence, for a color category {red, white, blue}, $n_{\text{color}} = 3$. Since each value is assigned its own coordinate, category i is represented as an n_i dimensional vector, y_i . Hence, the total number of dimensions used to describe the full set of n categories (y_1, \dots, y_n) is

$$\sum_{i=1}^n n_i$$

Note that sparse methods are especially useful here, since a categorical vector y_i will typically consist of mostly zeroes, with a single non-zero coordinate representing the categories' value (i.e., we encode the color red, using the previous example, as $(1,0,0)$). Note also that category vectors with different values are treated as orthogonal by the system.

A final issue is the representation of text. As described in previous related patents, all relevant blocks of text in the database are converted into a dictionary that maps unique strings to the number of times they appear in the database. An appropriate TFIDF weighting function is chosen and calculated for each of the p words that appear in the dictionary. The full set of text connected to a single customer can thus be represented as the vector (z_1, \dots, z_p) , where each z_i equals the number of times the word i appears in text related to the particular customer multiplied by the TFIDF score assigned to word i .

In summary, when a database describes its customers using a combination of numerical values, categories, and text, customer i can be represented by the vector $c_i = (x_1, \dots, x_m, y_1, \dots, y_m, z_1, \dots, z_p)$.

CONFIDENTIAL

6.62 An Example Profile Vector

Suppose we have a database containing information on customers' ages, their musical preferences (i.e. an answer to a survey asking: "Which do you prefer, Mozart or the Beatles?"), and the contents of the emails they've written. Furthermore, suppose the only salient variables in all the emails written consist of the words "Beatles", "Mozart", and "practice", and that we are using the function

$$TF / IDF(x) = \frac{1}{\sqrt{n_x}}$$

Where n_x represents the number of times word x appears in the dictionary. So, if the word "Beatles" appears a total of 217 times across the full set of customer emails, $n_{\text{Beatles}} = 217$, and

$$TF / IDF(\text{Beatles}) = \frac{1}{\sqrt{217}} = 0.067$$

We now want to represent one of the customers in the database; he's a 10-year-old boy who prefers Mozart to the Beatles, and who wrote an email to his friend that mostly describes his attempts at practicing Mozart, but in passing mentions his sister's new Beatles CD. Suppose he uses the word Mozart 2 times (although it appears 456 times in the full database of all customers' emails), the word Beatles 1 time (appears 217 times in database), and the word practice 3 times (appears 77 times in database).

We define the following coordinates:

$$x_1 = \text{age} = 10$$

$$y_1 = \{\text{Mozart, Beatles}\} = (1, 0)$$

$$z_1 = \# \text{ of times customer uses word "Beatles"} \times TF/IDF(\text{"Beatles"}) = 1 * 0.067 = 0.067$$

$$z_2 = \# \text{ of times customer uses word "Mozart"} \times TF/IDF(\text{"Mozart"}) = 2 * 0.047 = 0.094$$

$$z_3 = \# \text{ of times customer uses word "practice"} \times TF/IDF(\text{"practice"}) = 3 * 0.114 = 0.342$$

In our example, then, we might encode this boy as customer 1:

$$c_1 = (x_1, y_1, z_1, z_2, z_3) = (10, 1, 0, 0.067, 0.094, 0.342)$$

6.63 Choosing an Appropriate Level of Data Granularity

We define the term granularity to denote the level of detail available within a given set of data, which is often structured hierarchically. Suppose a grocery store database contains records for a box of flavored gelatin powder. This could be categorized in a variety of ways; moving from the most specific to the most general, we might treat this data point as "12.5 ounce, strawberry flavor, Jello-brand gelatin dessert" (which would be entirely different from "12.5 ounce, banana flavor, Jello-brand gelatin dessert"), or as "12.5 ounce Jello gelatin" (a categorization which would treat as identical the strawberry and banana Jellos), or as "flavored gelatin", or as "dessert", or as "food", or as "grocery".

When analysis is performed on such data, the level of granularity chosen will have a strong effect on the outcome of the analysis. If the level of granularity is too fine-grained, the data will be too sparse, although it could be potentially aggregated to the next highest level of granularity. If the granularity is too coarse, the results of the analysis might be overly general (e.g., a customer would find a collaborative filter useless if the only recommendation it makes for a dessert choice is "go to the grocery section of the store").

Since the level of granularity will have a salient effect on the outcome of an analysis, it should be chosen very carefully, and might well play a factor in pricing when a vendor chooses to sell its data.

6.64 Methods used for data analysis

CONFIDENTIAL

1 In order to perform a wide range of analytical tasks, SDI needs to make use of a variety of computational approaches.
2 These are described below, starting with the simplest methods first.

3 • (1). Standard Database Searches

4 Since most of the data will be stored in centralized databases, simple searches, queries, and data filters can be
5 implemented by means of standard SQL commands. Typically, data will be collected or sorted using efficient
6 database calls before being fed through analysis routines; once complete, the results can be fed back out to the
7 database environment for further efficient manipulation.

8 • (2) Metrics – Measuring the Similarity Between Profile Vectors

9 Given two customer (or vendor) profiles, c_i and c_j , it is frequently desirable to know how similar they are. For
10 this purpose, we define the similarity metric $M(c_i, c_j)$ to be a function that takes as input two customer vectors and
11 returns as output a numerical value in the range $[0,1]$. When two customers c_i and c_j are identical, $M(c_i, c_j)=1$;
12 when they're completely different, $M(c_i, c_j)=0$.

13 The problem is somewhat simplified by the fact that we treat all customers as vectors. Given two customer
14 vectors, we can use the correlation between them to serve as our metric:

15
$$M(A, B) = \cos \theta = \frac{A \cdot B}{\|A\| \cdot \|B\|}$$

16 Note that θ here represents the angle between the vectors A and B, and that we expect all coordinates of the vectors
17 to be positive (in order for $M(A,B)$ to keep its output in the range $[0,1]$).

18 In more complicated cases, however, a customer vector might contain multiple fields with varying ranges of
19 values. For example, we might have customer vectors of the form $c_i=(age_i, income_i)$, in which the maximum age is
20 80, but the maximum income is 300,000. In such cases, the coordinates with larger values will dominate the
21 similarity metric, overwhelming any influence that smaller fields might have.

22 This requires a normalization of the customer vectors, which can be done in several different ways. One
23 approach would be to scale every coordinate by the maximum observed value, forcing all coordinates to lie
24 between 0 and 1 (again, enforcing the rule that all coordinates must be positive).

25
$$c_i = \left(\frac{age_i}{\max(age)}, \frac{income_i}{\max(income)} \right)$$

26 The only problem with this is that if a coordinate's maximum value is an outlier (being vastly bigger than the
27 typical value), most of the coordinates' values will seem unusually small once they are scaled by the maximum. In
28 such cases, it might be better to scale the values with a "squashing" function such as the sigmoid, which deadens
29 the impact of extreme values; one such configuration would be the following:

CONFIDENTIAL

$$1 \quad \overline{age}_i = \frac{age_i - \text{mean}(age)}{\sigma_{age}}$$

$$2 \quad \overline{income}_i = \frac{income_i - \text{mean}(income)}{\sigma_{income}}$$

$$3 \quad c_i = \left(\frac{e^{\overline{age}_i}}{1 + e^{\overline{age}_i}}, \frac{e^{\overline{income}_i}}{1 + e^{\overline{income}_i}} \right)$$

4 Note that the mean and variance of the data points are used to fully normalize them, such that the sigmoid
5 function will spread the values somewhat more evenly between zero and one.

6 The previous approaches are especially useful for single numerical fields, which might well overwhelm each
7 other if some sort of normalization isn't performed.

8 A different problem arises for text or large categorical fields, since they can potentially consist of hundreds of
9 coordinates capable of overwhelming the influence of single numerical fields. Suppose we believe the age of a
10 customer is as important as the text of articles read. In such a situation, the thousands of coordinates devoted to the
11 text field would dominate the metric's behavior, negating any influence that age would have on our measure of
12 similarity – clearly not a good situation.

13 A solution to this would be to find the correlations among the fields taken separately, then average the result.
14 That is, if each customer $c_i = (age_i, \text{text}_i)$, where text_i is a vector with a very high number of dimensions, we could
15 define the metric:

$$16 \quad M(c_i, c_j) = \left(\frac{\text{corr}(age_i, age_j) + \text{corr}(\text{text}_i, \text{text}_j)}{2} \right)$$

17 Where

$$18 \quad \text{corr}(c_i, c_j) = \frac{c_i \cdot c_j}{\|c_i\| \cdot \|c_j\|}$$

19 The result is a metric that gives equal influence to each field.

20 • (3) Forming Vectors Into Groups

21 The process of classification is essential to collaborative filtering, as it allows different vectors to be formed
22 into groups based on some measure of similarity. If we are able to create groups of customer vectors, for example,
23 we can then give individual customers recommendations based on the patterns of their group-mates, who
24 presumably have similar tastes.

CONFIDENTIAL

K-means Clustering and Nearest Neighbor algorithms are extremely useful for grouping purposes: previous iReactor patents <<FRED - references?>> give a full and detailed description of our customized versions. This section gives a brief overview of these methods.

(3.1) Clustering

K-means Clustering is an algorithm used to partition a coordinate space such that all vectors in a given partition are more similar to that partition's vector average (the centroid), than to the centroids of any other partition. It is a process that iterates over the following steps:

0. "Seed" the coordinate space with the initial centroids, which are vectors used to describe the centers of the clusters, in the sense that they are the average of all the vectors currently assigned to the partition. This can be done randomly (assigning centroids random coordinates) if no other information is available, or it can be guided by pre-existing information. For example, if we wish to cluster vectors of music customers, we can use information about musical genres to create initial partitions that correspond to pop, gospel, classical, etc. This will locate the centroids in well-spaced intervals across the coordinate space.

1. Assign vectors to the most similar centroids. This is done for each vector by scanning across all centroids and calculating similarity $M(\text{vector}, \text{centroid}_i)$; once finished, the vector is assigned to the cluster whose centroid has the greatest similarity. In this stage, vectors may switch their allegiance from one centroid to another, if the relative distances to the vector have changed sufficiently since the previous iteration. If no vectors change their allegiance, the iteration process is complete, and the algorithm stops.

2. If the iteration is not complete, recalculate the centroids by setting them equal to the average of those vectors that have been assigned to them. Go back to step 1.

Once the algorithm converges, the vectors are grouped into clusters. The centroids' coordinates as well as the identity of cluster members is useful information that can be passed on to subsequent stages of analysis.

(3.2) Nearest Neighbor

The nearest neighbor algorithm, simply stated, creates a list of those vectors in a database that most resemble a particular target vector. This is accomplished by comparing the target vector, in turn, to every other vector in the database; the similarity between them is recorded, and once the comparison loop is complete the list of similarities is sorted. The top k members of this list are returned as representing those k vectors which most resemble the target.

(4) Generalizing Across Databases

One of the most useful aspects of SDI is that it allows for inferences to be drawn across different databases through underlying connections in membership or content. An especially strong link can be made between commercial databases if they have customers in common. However, for reasons of privacy, individual customers may choose to use different pseudonyms when dealing with different vendors. This might be preferred by the individuals, but it weakens the inferences that can be made between fields occurring in different databases.

The techniques chosen to infer correlations across different databases will depend on how many pseudonyms are shared in common. At one end of the spectrum, every customer uses a single pseudonym for all transactions,

CONFIDENTIAL

1 and makes an appearance in every database. At the opposite end of the spectrum, every customer uses a different
2 pseudonym with every vendor, and may appear in only a single database.

3 Case 1: All customers use a single pseudonym, and appear in all databases considered.

4 This is the simplest situation to handle. Since all customers appear in all the databases, the customer vectors'
5 fields are essentially scattered across several locations, but can be easily reconstructed. For each customer, we
6 define a new data vector that concatenates that customer's representation from across the different databases.

7 Hence, if we are considering databases A, B, ..., Z, and customer i appears in each one, we define a new
8 vector $c_i = (c_{Ai}, c_{Bi}, \dots, c_{Zi})$, where c_{Ai} is customer i's vector in database A. We then proceed as usual, making
9 inferences with these augmented customer vectors.

10 Case 2: Most customers use a unique pseudonym, and frequently appear in different databases.

11 In this situation, although we see some connections between the databases, many pseudonyms appear in only a
12 single location. Using Bayesian techniques, however, we can still make predictions for customer vectors across
13 databases.

14 Suppose we have a set of databases, A, B, ..., Z. Taking each database in turn, we cluster it using all available
15 data. Thus, using every record in database A, we group A's customers into clusters

16 A_1, A_2, \dots, A_n . Taking database B, we create clusters using all of B's information, creating customer clusters $B_1,$
17 B_2, \dots, B_m , and so forth.

18 Now, scan both databases for common pseudonyms (representing those customers who have interacted with
19 both vendors under the same pseudonym) and create count variables w_{ij} to represent the number of pseudonyms
20 that appear jointly in A_i and B_j .

$$21 \quad P(B_j | A_i) = \frac{P(B_i \wedge A_i)}{P(A_i)} \frac{w_{ij} / \text{total}}{\sum_{j=1}^m w_{ij} / \text{total}}$$

$$22 \quad \text{total} = \sum_{i=1}^n \sum_{j=1}^m w_{ij}$$

23 We can now produce the probability that a pseudonym appearing in A_i will appear in B_j :

24 For example, if we have a database of airline ticket purchases and a database of restaurant visits, we can create
25 clusters, in the first case, of customers who travel to similar destinations, and in the second case, of customers who
26 eat at similar restaurants. Given that a particular customer belongs to a cluster of people who frequent Caribbean
27 restaurants, we can infer which travel packages would most appeal to him based on the linking probabilities, as
28 defined above.

CONFIDENTIAL

• Multivariate Extensions:

- 2 If we have a third database C, and there are a large number of pseudonyms common to A, B, C, the above
- 3 probabilities can easily be extended. For example, knowing that a customer appears in A_i and B_j, we can calculate
- 4 the linking probabilities to any C_k:

$$5 \quad P(C_k | A_i \wedge B_j) = \frac{P(A_i \wedge B_j \wedge C_k)}{P(A_i \wedge B_j)} = \frac{w_{ijk} / total}{\sum_{k=1}^p w_{ijk} / total}$$

$$6 \quad total = \sum_{k=1}^p \sum_{i=1}^n \sum_{j=1}^m w_{ijk}$$

- 7 Or, if there aren't many pseudonyms that span all three databases, the probability of C_k given that a
- 8 pseudonym exists in A_i and B_j could be approximated by:

$$9 \quad P(C_k | A_i \wedge B_j) = P(C_k | A_i) \cdot P(C_k | B_j)$$

10 Case 3: All customers use several pseudonyms, and none appear in different databases

- 11 In this situation, there are no common customer codes that can be used to create links across the databases.
- 12 However, the mere fact that several databases have been brought together for analysis should imply that there are
- 13 semantic commonalities in the data.

- 14 Although each database contains different fields, it may be the case that those fields deal with related subjects.
- 15 A human expert, knowledgeable in the content of the databases, the subtleties of the domain, and the overall goal
- 16 of the analysis (e.g. the creation of recommendations), will be in a position to create a "common-information
- 17 profile" that spans the databases. In essence, the common-information profile defines a format that allows vectors
- 18 from different databases to share a common coordinate space.

- 19 The idea is this: the expert designs a high-level vector format that embodies the content deemed important for the
- 20 project goals. Next, for each database he develops a mapping that encodes the database's elements into the generic
- 21 format. Finally, the desired analysis is performed on the full set of common-information profiles.

- 22 Although the expert will have to create completely new fields for the common-information profile, certain types of
- 23 data will map directly to the common-information format. In particular, if every database contains text (catalogued
- 24 and counted, for TF/IDF purposes, by accompanying dictionaries), the union of the words will define the text
- 25 coordinates of the new common-information profile. When word counts are being mapped from their original
- 26 databases to the new vector, the original TF/IDF weightings may be used, or new TF/IDF weightings may be
- 27 created (using a dictionary constructed from all the databases' text taken together).

CONFIDENTIAL

- 1 Once analysis has been performed, certain common-information profiles will be grouped together by their shared
- 2 similarities, although the pseudonyms they represent may have been originally drawn from different databases.
- 3 Such groups will represent links between different databases, and may be used for predictive purposes (see end of
- 4 example).

5 6.65 Example of Cross-database Analysis

6 Suppose we have the following databases:

- 7 A. A travel agency keeps track of tickets sold, and vacation web pages browsed.
- 8 B. A bookstore keeps track of books sold, and stores an electronic version of the New York Times Review of Books.
- 9 C. A sporting-goods and clothing shop, keeps track of purchase items sold (which includes magazines, for which
- 10 electronic text exists).

11 A certain airline wants to promote various vacation packages it has available, which include both European

12 and Caribbean vacations, as well as singles and family packages. Although it has leased rights to databases A,B, and C,

13 it turns out that no customer pseudonyms appear in more than one database at a time – in other words, there are no

14 shared records.

15 A vacation expert is hired to create a common-information profile. He creates the following information

16 vector:

17 (list of tropical countries, list of European countries, family score, list of sports, text)'

18 Note that the family score is a numerical value ranging from 1 (young singles) to 10 (many small children),

19 and indicates what kind of person the customer is (a party-oriented student vs. a sedate father of three).

20 The expert the creates the following mappings:

- 21 A. Travel Agency. Link destinations of tickets sold to country fields (i.e., the number of trips to Germany by a
- 22 customer would be placed in the Germany field of the common-information profile). Link sales of children's tickets, or
- 23 requests for children's meals, to family score. Put web-page data into text field.
- 24 B. Bookstore. Link travel books' text to country lists. For all books purchased by a customer, map text from book
- 25 reviews into text field.
- 26 C. Sporting-Goods store. Map warm-weather clothing (and swim gear) to tropical countries, ski gear to countries with
- 27 skiing areas. Map sales of toys or children's clothing to high-value family scores, map revealing-bikini and student-
- 28 discount sales to low-value family scores. Map text from magazines purchased by a customer to text field.

29 These mappings are then applied to each database, generating a full set of common information profiles. These

30 are then clustered, forming groups that share commonalities.

31 The expert can now do several things with the results. First of all, he identifies the general "flavor" of each

32 cluster (e.g., families with small children that enjoy winter, Europe, and skiing); the pseudonyms contained within each

33 cluster can then be targeted for vacation packages suitable to their tastes. Secondly, the fact that pseudonyms from

34 different databases have been clustered together allows the expert to plan cross-category marketing. If certain travel-

35 book-buying parents have been grouped together with parents who bought their children swimsuits and scuba toys, it

36 may be that they share a preference for family activities that take place in warm places or by the seashore. Hence, the

37 book-buyers might be advertised various ocean-related sports goods appropriate for young families, and likewise the

38 swimsuit-buyers might enjoy getting recommendations for travel books that describe tropical destinations that are

39 especially fun for children. That is, if the goal is to cross-market items from A to customers in C, the most logical

40 source of recommendations would be the people in A who have been grouped with the people in C.

CONFIDENTIAL

- 1 SDI also allows dynamic and first-time personalization for a user that visits a vendor's site, on the basis of previous
2 analysis and the user's profile.

3 6.66 Methods for Validation

- 4 The main result of using SDI is the creation of connections between different data points, linking vendors and
5 customers, customers and recommendations. To a large degree, the overall success of an SDI analysis is the relevance
6 of the connections that are inferred from the data. It is often the case that a certain amount of validation is required to
7 determine which analytical approaches are the most successful, given that the analyst has had to choose a particular
8 combination from a wide range of algorithms, data sets, levels of granularity, and parameter settings. The process of
9 validation measures the relative success of a given project, and is used to guide the analyst through further iterations of
10 tuning and adjustment so as to optimize the final results of the analysis.

- 11 There are two general approaches, not necessarily mutually exclusive, to validation: the first is fairly
12 quantitative, the second relies more on human expertise and intuition.

13 • (1) Quantitative Approaches

14 1.1 Test Against a Validation Set

- 15 The goal of validation, in this context, is to measure how successfully SDI makes a prediction, most
16 commonly a recommendation. Before a recommendation system can be used commercially (when it is exposed to
17 actual customers), it is important to make sure that it is using the best possible combination of algorithms, input
18 data, and parameter settings (e.g. TF/IDF tuning). If several different combinations are under consideration, there
19 is a need to gauge the relative predictive accuracy of one approach over another. This can be accomplished by
20 holding out part of the data set, training the recommendation system on the remainder, then evaluating the strength
21 of the recommendations made for the hold-out set.

- 22 Suppose we are testing two possible settings for a system that recommends music. We make a copy of the
23 customer purchase records and remove a single purchase at random from each customer – this slightly reduced
24 copy will serve as our training set. We then allow the two rival systems to recommend musical albums for each
25 customer, based on the information in the training set alone. Typically, these recommendations will take the form
26 of a list of items with corresponding numbers that indicate the strength of each recommendation. The relative
27 performance of a set of recommendations can then be gauged by looping across each customer, noting whether or
28 not the system recommended the item that had been held out, and if so adding it to a running total. The system
29 with the highest total can thus be judged the most effective, since it most strongly recommended items that the
30 customers did, in fact, end up purchasing.

- 31 Because the result of this type of validation is a quantitative score, it is possible to automate the model
32 selection process. Given a set of analytical approaches (each with its own array of parameter settings), it is possible
33 to loop through the full parameter space (using a grid of evenly spaced numerical values, if needed, to reduce
34 dimensionality), computing a validation score at each iteration. Those combinations of algorithms and parameter
35 settings that demonstrate the best performance could be chosen as the top candidates for the final system
36 configuration, since they do the best job at predicting customer behaviors.

37 1.2 Dynamic Approach

- 38 The problem with the hold-out approach to validation is that it isn't dynamic, since it doesn't reflect the
39 impact that the recommendation system has on the customers once it is implemented, and may be based on data

CONFIDENTIAL

1 that doesn't contain current trends. After all, it is better to predict what the customer will buy rather than what the
2 customer has bought in the past.

3 A better approach is to run a controlled experiment against the actual customer base. First, the pool of customers is
4 split at random into different segments. Next, each approach under consideration is used exclusively to make
5 predictions for a given segment. Once the trial period is over, each system is given a score based on how valuable
6 its recommendations turned out to be (this could be measured by total sales generated, for example, or by the
7 number of times a customer made use of a recommendation).

8 • (2) Human Expert in the Loop

9 Although quantitative methods can automate the validation process to some degree, at the beginning of many
10 projects there is so much raw input data available and so many decisions that have to be made about the analytical
11 approach that an automated process would have to test a prohibitive number of combinations of data, algorithms,
12 and parameter settings to get optimal results. In such cases, it is useful to employ a human expert who understands
13 the psychology and nature of the particular domain being analyzed.

14 Such a person will have intuition about what is and isn't relevant for his domain. For example, a movie expert
15 might be called in to work on a movie-recommendation system, for which an immense amount of input data is
16 available. In choosing relevant fields for analysis, the expert's understanding of cinema would lead him to include
17 the director's name and numbers of Oscars awarded, whereas the exact length (in minutes) of the movie would be,
18 in his estimation, irrelevant and therefore excluded.

19 Once the analysis is complete and recommendations have been made, the expert's opinion (based on a
20 qualitative understanding of the domain) can be used to guide which particular combination of settings, chosen
21 from a list of candidates with detailed test outputs, should be used for the recommendation system.

22 • (3) Combined Use

23 There is certainly no reason why both approaches couldn't be used in combination. Many data sets include fields
24 that are extremely noisy or simply irrelevant to a given problem; a human expert can be employed to pare the data
25 set down to a reasonable size and dimensionality, using his domain expertise to create a data model reasonable for
26 the proposed analysis. Next, automated methods can be used to fine-tune the parameter settings and to choose
27 which subsets of the input data are the most useful. Finally, the human analysts called back to qualitatively
28 evaluate the results of the fine-tuning, making the decision to either start a new iteration of the analysis, or to
29 certify that the process is complete and ready for commercial application.

30 6.7 Conditions on Usage and Privacy

31 As previously described in this patent, various conditions can be placed on the way in which a set of data may be used
32 (i.e., can the user make a personal copy of the dataset?), as well as on the privacy controls put in place. It might well be
33 that a vendor is willing to share only a portion of his database, or that he will release only randomized aggregates in
34 accordance with the level of privacy he has guaranteed his customers.

35 Although such restrictions could impact the content of the data analyzed by a vendor, as long as it is kept in an SDI-
36 compliant format it can be analyzed by SDI's suite of tools.

37 However, the data that is stored in the central SDI sever still has tight usage restrictions, as placed on the data by the
38 submitting vendor and the user that the data pertains to. For example, the user will have specified a use-of-data policy
39 that could restrict the data to be used for only personalization purposes, or only solicitation purposes. In this case the
40 user has connected to the site of a vendor, so this use-of-data restriction is irrelevant. The vendor will place additional
41 restrictions on the data, for example the vendor might wish that the data that it submits is never used to personalize the

CONFIDENTIAL

1 service offered by a direct competitor, with the same classifier label. When this is the case the central SDI server must
2 not release any of this data to the vendor, even if this vendor is not a direct competitor, because there can be no
3 guarantees that the vendor will not pass the data on to a direct competitor.

4 The data within the central SDI server is therefore in two main classes, data that can be released by SDI because the
5 vendors that have submitted the data have placed no restrictions on the other vendors that can access the data, and data
6 that can never be released by SDI because the vendors that submitted the data have placed restrictions on the other
7 vendors that can access the data. We can still make use of the restricted data when providing profiling information to
8 non-competitor vendors. Periodically the central SDI server performs collaborative filtering for each vendor, on the
9 pseudonym records within its user base that contain information that relates to the business of that vendor.

10 For example, consider a vendor that sells compact discs. The vendor will have submitted profiling information about
11 the each user (represented with a pseudonym) that it has done business with. Furthermore, the central SDI server might
12 also contain additional profiling information that relates to the same pseudonyms, that has been submitted by other
13 (non-competitor vendors). Finally, there may be records about other pseudonyms that have purchased music - either in
14 compact disk format or an alternative format - that has been submitted by vendors that are happy to have the
15 information used to help directly competing vendors.

16 The central SDI server can provide the vendor with the benefit of all of this data without actually releasing data to any
17 of the vendors, by providing the results of the collaborative filtering analysis, restricted to the vendor's own data
18 model. For example, given the connection certificate, that certifies that a user with pseudonym P has just connected to
19 the site, the central SDI server can look up the profile of the pseudonym, and make appropriate product
20 recommendations to the vendor - recommendations based on the analysis that has been performed using all of the data
21 that it has available.

22 The structure of a typical data record contains a number of additional fields to represent profile-usage policies.
23 (Pseudonym, data, L, R, P, S). The 'pseudonym' is the public key and ISP-level proxy server IP address relating to the
24 pseudonym, the 'data' field is the profile information for the pseudonym - in general a sparse vector of numbers, L
25 contains vendor classifiers that are excluded from the data, R is a {0,1} bit that indicates whether any of the data is
26 randomized, P is a {0,1} bit that indicates whether the data can be used for personalization of service to the user when
27 he/she visits a site under the pseudonym, and S is a {0,1} bit that indicates whether the data can be used for
28 solicitations to the user with the pseudonym.

29 6.8 Location of Data and Algorithms

30 Although SDI might have controlled links to vendors' databases, the actual information might not physically reside
31 within the SDI system. One could imagine a vendor requiring the joint analysis of data that includes highly proprietary
32 information (kept for safety behind the company's firewall) with slightly less-critical information belonging to another
33 vendor that is stored by SDI. In this situation, the vendor would have the ability, as with a lending library, to "check
34 out" both the secondary data and relevant algorithms from SDI, and to use them at their own site. Thus, although both
35 data and algorithms might reside either at vendors' home locations or within the SDI system itself, the general analysis
36 will work transparently across these boundaries.

37 6.9 Other Issues

38 For security reasons, the contents of databases may be injected with a small amount of noise. This prevents database
39 users from surreptitiously connecting database records to individual customers, yet maintains the quality of inferences
40 made about the database in general.

41 Although such "noisy" records don't pose too much of a problem for those methods that make generalized inferences,
42 it should be noted that recommendations made for individual customer vectors that have undergone such randomization
43 will be less useful, since predictions are being made for a noisy target.

CONFIDENTIAL

- 1 A final consideration is the reduction of the data vectors' dimensionality (which can be extremely high), since it is
2 harder to make clean inferences about sparse data. There are many standard methods that can be used to achieve this,
3 such as Principal Components Analysis.
- 4 Another approach is to adjust the granularity of the data, if at all possible. In a music store analysis, for example, there
5 might be many more album titles that artists (since each artist can produce multiple albums). In such a case, purchases
6 could be recorded by artist rather than by album, greatly reducing the dimension of the customer vectors' purchase
7 space.

3 7. *Randomized Aggregates: Enhancing User Privacy*

- 9 Randomized aggregates provide a cheaper and more secure alternative to cryptographic techniques, such as secure-
10 function-evaluation for providing information without compromising privacy. Even with pseudonyms it is possible
11 that the identity of an individual can be revealed through revealing too much specific information about the
12 transactions/profile of an individual. This occurs when the information places a strong restriction on the set of users
13 that could satisfy the revealed constraints. We enable vendors and users to report information about pseudonymized
14 users in a randomized form that prevents this type of reverse-engineering to identify the user behind a pseudonym. We
15 add randomization to data fields so that the user's privacy is protected, but also so that the data still allows accurate
16 aggregation and collaborative filtering/multi-attribute clustering analysis. Randomized data is also secure to
17 computational attacks and the loss or theft of private keys-because we degrade the data, and make access to any one
18 data item virtually useless.

- 19 This technique of randomized aggregates enhances user privacy guarantees, allows vendors to disclose useful
20 information without violating user-privacy requests, servers to reduce system-wide reliance on cryptographic solutions.
21 The basic idea is to add a small noise term to each field of a user's profile. Aggregate data can allow economic
22 evaluation of data without access to the data, although aggregate data can have some inherent value in itself.
23 Randomized aggregates enable: (1) personalization; (2) aggregate statistics; (3) protection against profiling.

- 24 In many situations it is sufficient to gather information about the activities of an entire demographic group rather than
25 about any one individual. For example, a VCR dealer might be interested in the chance that a person buying a new
26 television set will purchase a VCR within the next twelve months. However, information about whether any one
27 individual who purchases a TV also goes on to purchase a VCR within a year should not need to be revealed in the
28 process of computing this chance.

- 29 Traditional cryptographic methods are capable of solving the problem of revealing only aggregate information while
30 concealing individual information. Methods exist for computing aggregates or other values from encrypted information
31 without first decrypting this information. (Such methods are dealt with in an area of cryptography known as secure
32 function evaluation.) However, the general-purpose nature of these methods makes them unnecessarily cumbersome for
33 the limited problem at hand. In particular, the communication and computation requirements of these methods when
34 applied to the problem of aggregation result in an unacceptable overhead on the system.

- 35 There is also an additional problem with such cryptographic techniques. Secret information can be compromised by
36 successful computational attacks on the cryptographic scheme or by the loss/theft of private keys. Such problems are
37 present in all uses of cryptography. Nevertheless cryptography is used where it is the best alternative. However, once
38 again, for the limited problem at hand our solution is safe both from computational attacks as well as from private keys
39 being compromised.

- 40 The above-described problems are solved by the use of a new technique that randomizes sensitive data. User decisions
41 and other numeric fields of records are modified by the addition of a "noise" term that may be positive or negative.
42 These noise terms are chosen by sampling from carefully chosen probability distributions. Different, independently
43 chosen noise terms are used for each field of the record that needs to be perturbed by the addition of noise. The
44 modified values of these numeric fields are then transmitted to an aggregating database. The net result is that when very
45 few records have been aggregated the value of the numeric fields are likely to be quite unreliable. However the
46 accuracy of these fields improves gradually as more and more records are accumulated.

- 47 As an example, consider the following situation. Suppose that we want to determine the average income level of people
48 purchasing camcorders. In each record there could be a numeric field indicating the type of object that we are dealing

CONFIDENTIAL

1 with. For example, assume that this field is "I" exactly when the object in question is a camcorder. Now suppose that
2 each of these records has another field that contains income information. If this information has been perturbed by
3 random noise then any individual record in the bucket has income information that is highly unreliable. However, when
4 sufficiently many records are aggregated the income information becomes more accurate! (This is a consequence of a
5 theorem in probability theory called the Law of Large Numbers.)

6 As another unrealistic but illuminating example, consider the following scenario: Suppose an organization receives
7 monetary contributions from a number of individuals, who for reasons of privacy, wish to conceal the amount of their
8 contributions. Suppose that the total of all contributions is to be public knowledge. The randomization scheme works
9 as follows: When an individual contributes x dollars, she chooses a random number r according to a specified
10 probability distribution (with mean 0) and sends $x+r$ to the aggregate database. By allowing widely dispersed values of
11 r , we can ensure that an eavesdropper who obtains the value of $x+r$ has very little information about x . However, if a
12 large number of individuals register their contributions in this manner, then the total of the values sent to the aggregate
13 database will be very close to the true total contribution. The actual scheme and its variations are more complicated as
14 they have to deal with multiple fields and more sophisticated attacks on privacy. But the basic idea described in this
15 example will be used repeatedly in these schemes. In this example, the total of a set of values was required to be public.

16 7.1 Adding Noise to Fields

17 A record is a tuple of information containing various fields some of which are numeric, for example, describing the
18 details of a commercial transaction. Noise or perturbation which refers to the random value added to individual numeric
19 fields of records.

20 If the field is a continuous value, such as salary, then we can add a Normally distributed term with zero mean and a
21 carefully selected standard deviation. The standard deviation is "tuned" to provide a good tradeoff between individual
22 privacy and accuracy of aggregate analysis. For example, if an independent noise term is added to the salary field of a
23 set of user profiles, and a vendor requests analysis of the mean salary over the set, then with a large noise term more
24 users are required in order to generate an accurate average salary. In contrast, if only a small noise term is added to
25 each salary field, and a third party with knowledge of the salaries of users requested data on each user, the third party
26 would be able to match the users to the data fairly accurately, on the basis of the salary field-particularly for users with
27 distinctive salaries (very low or very high).

28 If the field is one of a discrete set of items, such as the name of a CD that a user has purchased, then randomization
29 proceeds slightly differently. In this case, one replaces the identifier with an identifier that is semantically close. For
30 example, another CD of the same genre.

31 We can still perform correlation across fields with randomization, so long as the randomization does not destroy any
32 trends between fields. Randomized data is marked as such within SDI, and labeled with the degree of degradation, so
33 that SDI can be aware of the number of records to get relevant accuracy levels, and can report accuracy to customers.

34 We need to add noise to make data elements "close" to the accurate values. With discrete data, such as the name of an
35 artist, "close" must be defined within the correct metric. The appropriate metric is such that a "close" value shares
36 many of the same characteristics. For example, it is not appropriate to assign a close value on the basis of a shared last
37 letter in the first name, but it is appropriate to assign a close value on the basis of an artist from the same genre of
38 music. We call such clusters of artist names "semantic clusters", to imply that they have meaning in the domain.
39 Semantic clustering that enables useful randomization of discrete field can be automated when goods are frequent
40 purchase, high volume goods-where individuals purchase goods on multiple occasions, and more than one of the family
41 of goods on a single occasion. High price, low volume goods, should be randomized on the basis of expert analysis (for
42 example new cars, computers...) - where an expert can extract key features of a purchase, and represent the purchase
43 generically using either a single prototype good, or one of a set of approximately equivalent goods.

44 We do not require that every vendor/user uses the same distribution for its additive noise term. The choice can be made
45 autonomously. All that is required is that a vendor specifies the degree of randomization, so that SDI can be aware of
46 the type of data that it is aggregating and selling.

47 We now begin an elaboration of the variations possible to the basic scheme and the assumptions and requirements
48 under which each variation would be appropriate.

CONFIDENTIAL

1 Records are associated with user pseudonyms, not user IDs. All fields can be perturbed with noise, except the
2 pseudonym ID field, and any database query can be transformed into a query on "fuzzy" fields. Queries are
3 dynamically constructed, and it is necessary to be sure that the result set contains enough data points so that a user's
4 privacy is not compromised. Furthermore, we do not normally allow restrictions of pseudonym IDs.

5 The system can be configured and used in a number of different ways. Each transaction involves one or more parties.
6 Most common transactions involve two parties, the vendor and the customer. It is assumed that the communication
7 between these two parties involving the actual transaction is made secure by the use of a public-key cryptosystem or
8 otherwise. This assumption will be valid in most electronic commerce schemes. Once the transaction is completed, the
9 vendor or the customer is responsible for transmitting a record to the central SDI server. Again it is assumed that this
10 communication is secure.

11 There are several mathematical issues to be resolved, foremost of which is the choice of probability distribution for
12 noise. The system proposed here will be "tunable" for each application and it will be the application that will determine
13 the appropriate noise distribution. Here we simply describe the issues involved in the choice of the noise distribution.
14 An important requirement on the noise distribution is that it should have expected value 0 since we want the noisy total
15 to converge to the true total as more and more entries are aggregated. The second consideration is the variance or
16 standard deviation of this distribution. If the variance is made large (relative to the actual values involved) then each
17 noisy value reveals little about the true value of a field in a record. Clearly this is desirable in order to preserve privacy.
18 The drawback is that when the variance of the noise distribution is large, a large number of entries have to be
19 aggregated before the sum of the noisy values approaches the sum of the true values. Thus there is a tradeoff between
20 the level of privacy protection and the level of aggregation at which responses to queries become accurate. Once again,
21 each application can determine the point on this trade-off curve most suitable for that application.

22 The choice of the degree of noise to add to a data field is a tradeoff between protecting the privacy of an individual and
23 allowing robust profiling in the aggregate, and personalization at a pseudonym level. There is some information that
24 can be gained from a "fuzzed" numeric value. Suppose that the noise terms are added from a Gaussian distribution, of
25 unknown variance and zero mean. A third party can gain information on the likelihood that a fuzzed zero-one variable
26 has a real value of one as follows: monitor a stream of fuzzed values and fit the most likely Gaussian noise distribution,
27 or if the data is present in a database, just fit the most likely additive noise distribution. Then, given an individual
28 fuzzed value x' and true value x , and an estimate of the statistics of the additive Gaussian noise distribution, it is trivial
29 to compute the $\Pr(X = 1 \mid X' = x')$. For example, if the noise distribution is distributed with mean 0 and variance 2, and
30 the fuzzed value, x' , is greater than 3, then it is more likely that the underlying true value, x , is 1 and not 0.

31 This is less of a problem when (a) the domain of the numeric attribute is large and/or (b) the variance of the noise term
32 is large with respect to the domain and/or variance of the numeric attribute.

33 We can use a cryptographic technique to verify the distribution of noise that is added to data - and also to enable replay.
34 A vendor must keep a record of the non-randomized data that is supplied. When generating a random perturbation, the
35 vendor uses a one-way function f on the object X to generate a seed for a pseudo-random number generator. The
36 pseudo-random number generator then generates a sequence of random numbers that are used to create the random
37 perturbation from a well-defined algorithm. SDI can use this technique to verify randomization, and audit a profile-
38 updating agent. "Playback ability" - the ability to reconstruct the original record from a noisy version of that record is
39 important for a number of purposes. An individual may want to obtain proof of a transaction for legal purposes and law
40 enforcement agencies with appropriate warrants might want to examine original records. Each client-level SDI proxy S
41 maintains a trapdoor function f (such as the RSA encryption-decryption function). When adding noise to a record the
42 proxy uses the fixed fields of the record as argument x and computes the inverse off when applied to x . (Note that
43 third parties that do not have access to the trapdoor secret will not be able to compute this inverse function.) The proxy
44 then uses this value as the seed for a pseudo-random number generator and uses the bits produced by the generator as
45 the random bits used to produce the noise. With this scheme it is possible for the agent database to "playback" the
46 noise perturbation and produce the original record from the noisy record. This playback scheme is optional and may be
47 used in an application if the feature is desired.

7.2 Randomized Aggregates: A cheap alternative to Secure Function Evaluation

49 Randomized data also enables SDI to release data to third parties for analysis, but still apply its results to non-
50 randomized data. The third parties cannot make use of data because they do not have the pseudonyms, and they cannot

CONFIDENTIAL

1 identify any of the users. The models that they generate can be sold to SDI, matched with pseudonyms, and then rented
2 to vendors.

3 There are two uses of randomized aggregates:

- 4 (a) Revealing aggregated data from a database does not reveal information that is specific to a particular user, and
5 does not reveal information that is detailed enough to be useful to a vendor that wants to perform target
6 advertising. It is information of the form "if you had the details, this information would be useful", and allows
7 some calculation of probable economic value.
- 8 (b) When only aggregated data is required, then data can be transmitted securely in "fuzzed" form, and the aggregate
9 date reported accurately given enough records. This is useful when the owner of the aggregated database is not
10 trusted, or communication channels are not secure.

11 We need to check after a query that the data revealed does not compromise a user's privacy. A check at the time of
12 submission of new data to the database is not sufficient because: (a) initially, no data is secure, even in randomized
13 form - we can avoid a bootstrapping problem by checking the results of queries; (2) there are an infinite number of
14 queries that one would need to anticipate at the time that data is checked in. We can use a statistical test to establish the
15 minimum number of records that will ensure safe revelation of aggregate information, based on statistics/simulation
16 from a real set of data, or statistics over a typical user population.

17 We would like vendors that hold useful information to be able to trade that information. This requires establishing that
18 the information will be useful without revealing it. The classic solution to this would be Secure Function Evaluation.
19 Secure function evaluation of the expected economic value of receiving information from one database to enhance the
20 information present in another database is computationally expensive. Not only must the inputs for the function to be
21 evaluated be transmitted over the network, that is the full representation profiles, but the functions must also be
22 evaluated cryptographically.

23 An alternative to full secure function evaluation would be for an interested party to request aggregate information.
24 Ideally this aggregate information would allow the economic advantage of the precise details of the new information to
25 be estimated, without allowing the economic advantage to be immediately realized. The problem with using raw
26 aggregated information, such as what fraction of users in your database that bought a TV also bought a VCR within a
27 year, is that the requestor might already have information for many of the same users. What is required is that the
28 requestor also transmits a constraint on the user ids that he/she are interested in.

29 Here is a simple scheme for computing the economic value of new information, for Vendor 1 that is interested in
30 acquiring data from Vendor 2:

- 31 i) Vendor 1 sends Vendor 2 Query: What is the average income of the users in your database that have purchased a
32 TV in the past 12 months, but are not in this list of users. (Vendor 1 sends Vendor 2 a list of users that it already has
33 information about for this field).
- 34 ii) Vendor 2 can then send the aggregate information back.
- 35 iii) Should Vendor 1 be interested in receiving some form of access to information about the aforementioned set of
36 users, then negotiation over terms of contract can proceed.
- 37 iv) If the negotiation is successful, then finally Vendor 2 can send Vendor 1 the information, authenticated in some
38 way, and Vendor 1 can check the accuracy of the aggregate information that was reported in step (ii).
- 39 v) Payment for the information is made contingent on the accuracy of the information provided in step (ii).

40 Another interaction could look like this:

- 41 i) Vendor 1 says "Do you have attribute X for these users". <sends list of user ids>
- 42 ii) Vendor 2 reports "Yes, for 89% of them".
- 43 iii) Vendor 1 makes a deal with Vendor 2, and they agree payment contingent that the data that Vendor 1 receives is
44 consistent with the aggregate information provided in step (ii), and verifiable.

45 Finally, third party certification is necessary to guarantee that Vendor 2 is providing truthful data for the unknown fields
46 of the records requested by Vendor 1. This aggregate data approach will be useful when aggregate data is sufficient to
47 perform an economic analysis on the utility of new information.

48 Aggregate information does have value. Consider a vendor that has performed some analysis on sparse data from
49 his/her own database and is seeking to support his/her conclusions by performing a similar analysis on the sparse data

CONFIDENTIAL

1 of another vendor. This analysis could proceed with aggregate information alone, and have value in confirming trends
2 observed in the vendor's own data. The vendor could use the analysis of data from a similar vendor to draw
3 implications about his/her own database of user information. SDI must perform secure evaluation on information in
4 this case. The two different methodologies for the determining the value of information prior to disclosure of that
5 information are:

6 Aggregate information has no direct value to a requesting vendor, but can be used to assess whether the raw
7 information will be useful.

8 In this scenario it would seem preferable to avoid the complexity of requiring a trusted third party to perform the
9 evaluation. Instead the aggregate information can be provided directly to a requesting vendor, allowing the vendor to
10 perform his/her own analysis. Aggregated information is always sufficient, because the only way that a vendor could
11 use non-aggregated but randomized information would be to aggregate the information himself. We would never need
12 to transmit randomized records to the requesting vendor, because randomized information is only useful in its
13 aggregated form anyway. For example, the vendor might request aggregate information on the users for which it
14 already has information within its own database, or aggregate information only for new customers. Should the
15 aggregate analysis by the vendor indicate that the information held by the disclosing vendor is valuable, then mutually
16 beneficial terms of disclosure can be agreed between the requesting vendor and the disclosing vendor, and the
17 disclosing vendor can send the requesting vendor that information in a non-randomized form.

18 Aggregate information has direct value to a requesting vendor, even without the raw data from which it is derived (as
19 described above).

20 In this case we will require a trusted third party to evaluate the potential value of information held by Vendor B for
21 Vendor A. This will require the specification by vendor A of constraints, and possibly a sophisticated algorithm by
22 which to evaluate the utility of the data in the aggregating database. Note that the variables used within such an
23 evaluation function must be limited to aggregate information (means and sums) because that is the only accurate
24 information that can be derived by the trusted third party from the aggregating database.

25 The trusted third party responds to the requesting vendor, and the requesting and disclosing vendors are then free to
26 pursue negotiation in order to establish mutually agreeable terms of disclosure. If the data that is finally requested by
27 the requesting vendor is aggregate data, then that can be disclosed directly by the third party. Otherwise disclosure must
28 be between the vendors themselves.

29 7.3 Randomized Aggregates: Further Uses

30 Mutually agreeable terms of disclosure of that data can be negotiated between the requesting and disclosing vendors. If
31 the data is aggregate form, it is disclosed directly by the third party vendor otherwise, the non-randomized data is
32 disclosed by the appropriate vendor to the requesting vendor. The aggregate database may be used for any of the
33 following purposes:

- 34 (1) The use of explicit queries submitted by a given vendor to determine whether another vendor (whose aggregated
35 data resides in the secure database) possesses data which matches the criteria of the external request (typically in
36 conjunction with certain stated constraints, as discussed below).
- 37 (2) A comparison evaluation of the vendor's aggregated data to determine whether and if so, which portions of the
38 database (possessed by which vendors) contain data which is of potential relevance to the vendor's database.
- 39 (3) The generation of "virtual customer lists" which are determined to be of value to the vendor (through the
40 evaluation process of either #1 or #2 above). The pseudonym proxy server may then be used to target the desired
41 users matching the requested criteria of the vendor based upon terms agreeable to the requestor and data providing
42 vendor based upon a per-impression or per-transaction model (through use of a changeable pseudonym).
- 43 (4) If sophisticated formulae and/or esoteric algorithms are used where evaluation/processing of a significant quantity
44 of a given user's data is required, it may be preferable to confidentially convey these formulae and algorithms to
45 the third party. The outputs from pseudonymous users or the list of users which achieve certain predefined outputs
46 may then be conveyed to the vendor.
- 47 (5) Disclosure of randomized data by the randomized database to a vendor in order to develop formulae, algorithms
48 and models.

CONFIDENTIAL

As discussed further below, the value of the data to the vendor may be estimated by:

- 1) (if the prospective disclosed data is users), the number of users matching the relevancy criteria.
- (2) (if the prospective disclosed data is statistical in nature), the degree of statistical confidence associated with it.
- (3) (if available), the effective measurable degree of benefit resulting from previously disclosed similar data (if applicable, with similar statistical confidence) as to that of the present data to be disclosed.

We can also determine the number of other "undifferentiably similar" users that the user's identity can be masked by such that they are unidentifiable from the user. This is easily estimated by comparing the randomized data set to a set of real data, developing a randomized version thereof and observing the degree of confidence in being able to predict the true identity of the randomized user profile of a given "user" (from the randomized data) as confirmed by the actual data set. The present system may also notify the user if/when the number of identical users falls below a minimum predefined threshold to ensure this requirement

8. Techniques of Secure Function Evaluation

While randomized aggregates address the user privacy concerns associated with enabling vendors to gain full access to user profiles for modeling and in order to develop algorithms that fully leverage the new information that can be collected on-line, there is another potential user privacy concern. In order to actually deliver ** personalized service ** to an individual user, it may not be sufficient to use that individual's randomized profile, or an individual might be reluctant to reveal the randomized profile that is associated with her pseudonym, if it does indeed carry useful information.

A technical solution to this problem is offered through secure function evaluation, where two parties can evaluate a function based on distributed information without gaining any information about the inputs other than what is revealed by the outcome of the evaluation. The result of this method is the secure automatic evaluation of the user's profile using the algorithms and tools originally derived from randomized aggregates. The output is assured to be accurate by guaranteeing that both the user data and the vendor's imported function are secured to tampering by either party. The client-level SDI proxy server maintains a user's accurate profile data, while submitting randomized profile information to the central SDI server.

In one variation, the secure function evaluation can be used to determine whether there is data in the databases of other vendors that will be useful to enhance the user and object profiles of a vendor. If another vendor's data can be enhanced by the type of data on the vendor's site (and amenable terms are agreed upon), the vendor can then provide the useful portion of his/her data. Secure function evaluation allows the evaluation of data without access to the data, so that a fair contract can be agreed between the parties. The evaluation and exchange (and associated negotiation) of the data, occurs on a trusted server or the SDI proxy server where the relevant data resides. Because this analysis, data collection, compilation and modeling procedure occurs securely and under the control of the third party operator or intermediary operating the distributed trusted server network (e.g. SDI) the integrity of the data as well as secure enforcement of the privacy policies of the parties which possess access privileges to use that data are assured.

Practically speaking however, some of the user profile data (e.g. direct mail orders) will be directly collected by the vendor (this is inevitably true for any purchase of physical products whether directly using a store card or mail order with the exception of anonymous physical mail described below) and/or some users may not require that their data be protected by a pseudonymous proxy server. Accordingly, the secure function evaluation may also be used on behalf of the vendor to assess what portions of the data on the user's database are of potential value to the vendor. This potentially could involve such database matching activities as a simple direct comparison of fields of data between the user and vendor database in order to determine missing entries and/or fields on the vendor database to more sophisticated inference of data fields and/or leveraging additional relevant statistics for supplementing a sparse matrix for purposes of predictive user modeling purposes. This provides the user with a high degree of privacy, as no portion of the user profile needs to be revealed to the vendor.

The vendor criteria for selecting both users and appropriate targeted ads (based upon their user profiles) from a profile database can be performed entirely autonomously by the use of the present secure function evaluation method. Once the secure function evaluation has identified and qualified the relevant data and also measured the predicted benefits to the vendor, the vendor may possibly choose to negotiate terms for full purchase of the user data which is specifically relevant to that vendor. This data can be down-loaded to the vendor's site to directly analyze his/her user behavior

CONFIDENTIAL

1 statistics. The requesting vendor's "representation profile" consists of the complete collection of target object profiles
2 to be refined.

3 The first step is searching for and identifying data within remote vendor databases which are likely to be the most ideal
4 candidates for cross advertising with that vendor. Secondly, the database(s) which contain valuable data are evaluated
5 for the estimated degree of potential commercial benefit the data could provide. Thirdly, this value estimation and the
6 data privacy policies of the disclosing vendor (based upon his/her interest in providing data to the requesting vendor)
7 form the basis for terms, if any, for the data transaction).

8 One example of two different types of synergistic data bases include vendor databases and a distributed user database
9 (e.g. on the proxy server). Large customer interaction and transaction object profiles on the vendor's site are likely to
10 be much more robust than those collected by a user side profiling system distributed over a multi-server network, or
11 even multiple networks. From the vendors perspective the breadth of domains containing products that can be
12 potentially cross correlated from data on the user's database is very valuable. Vendors can hope to reach customers
13 who have visited other similar sites and sites which tend to be visited by similar users. Thus, the ad network is able to
14 utilize the user database (containing virtually all of the data about what sites and target objects a user accesses) in order
15 to effectively identify other vendors whose products and content are most frequently accessed by the most similar
16 users (and types thereof), suggesting the best target customers for each other's advertisements.

17 The system can contact the remote vendor requesting secure access to his/her data through secure function evaluation
18 in order to determine the estimated degree of statistical benefit. The vendor may choose to disclose data in randomized
19 form in order to enable the ad network to improve its data model while avoiding disclosure of data for purposes of
20 targeting his/her customers. If the vendor lacks installed cluster code in his database, the system automatically installs
21 the code as part of the secure function evaluation. In order to save overhead of re-installation the next time that a
22 request is made to access the vendors database or to update the present model the system requests the user for
23 permanent installation of the secure function evaluation and associated analysis code. This enables the automatic
24 updating of a distributed cluster model in accordance with dynamically changing information and user behavior
25 patterns. The cluster code is actually loaded onto the remote site, the analysis by the secure function evaluation may
26 also determine the estimated potential benefit achievable to the remote vendor (or user) associated with accepting the
27 request for his/her data. This may be achieved by comparing previous similar situations in which disclosure of data
28 occurred, i.e. the marginal measurable commercial benefits resulting from previous data sets with similar degrees of
29 statistical sparseness including the quantity of pre-existing and newly introduced data within a similar domain or cluster
30 to the present one. The degree of estimated commercial benefit resulting from disclosure of the requested data may be
31 then calculated. A separate estimation is then made if the remote vendor agrees to disclose his/her data (in non-
32 randomized form) for purposes of targeting ads to customers on his/her site (or email list).

33 Depending upon the existing degree of statistical confidence which is estimated between each target customer of the
34 remote vendor and the target objects of the advertising vendor, for the most part, it is likely that utilizing techniques
35 which improve the statistical confidence between users and target objects of these metrically similar vendors will yield
36 significant commercial benefits by improving the relevancy of the matches. It is also possible for a data privacy policy
37 to be submitted in conjunction with the release of any data. The vendor may desire that certain conditions of release be
38 associated with all or each portion of the data, in the form of explicit restrictions of usage which may be assigned and
39 tagged upon its release by the vendor. Adherence to these restrictions is the responsibility of SDI.

40 These data privacy policy restrictions may indicate which other vendors/vendor types (if any) may receive the
41 advertising benefits resulting from that data, what portions of that data can be utilized for the benefit of each
42 vendor/type. Also the pricing structure may vary in accordance with the type of data released, and who the vendors are
43 who will ultimately benefit from its incorporation. Additionally, the usage of the data may also be subject to controls
44 set forth by the contributing vendor. For example, applying different restrictions to different portions of his/her data,
45 restricting its modeling benefits to his/her own company exclusively, limiting the benefits to non-competitors or certain
46 vendors, limiting its usage to only certain advertising types (or purposes), restricting its usage to individual
47 impressions or sales (on a per impression or per transaction basis). Data may only be used on a per-impression or per-
48 transaction basis.

49 Because of the restrictions which vendors may place upon the usage (what portion of their data, to whom and/or for
50 what purpose), re-clustering the data model separately according to various degrees of disclosed statistical data can be
51 computationally expensive. We can use the same cluster model for portions of the model with common disclosure
52 policies. Each portion of the vendor's resulting data model with the same disclosure policy is called a "common
53 disclosure unit" or "CDU". Delineating CDUs within the data model involves continuous monitoring and identifying

CONFIDENTIAL

1 those portions of the data model (attributes and associated metrics). The CDUs, whether common access, limited
2 common access or totally "private", operate seamlessly with CDUs that are subject to less restrictive constraints.

3 It is also worthwhile to mention that more simple merging techniques may be used for enhancing remote marketing
4 databases by adding or filling in data fields to a sparse data model. This may be explicit data from matching fields, or
5 inferred data, e.g. additional products or attributes which statistically correlate with certain other products or attributes
6 or other user attributes. The system can notify vendors of the presence of potentially useful data within the user
7 database.

8 Secure function evaluation can also be used to retrieve target objects from secure private databases (or Web sites) that
9 would otherwise be inaccessible. SDI can identify whether there are any metrically similar target objects in the remote
10 database that match the target objects currently retrieved.

11 We use the cryptographic technique of secure function evaluation to evaluate the encrypted profiles of users that
12 pseudonymously access a vendor's site, but prevent the vendors from accessing the profiles directly. The profile-
13 analysis package is tuned for an individual vendor by analysis of the vendor's data model in SDI. The analysis is
14 provided to the vendor on a restricted basis by requiring that (a) the proxy servers request a new encrypted profile for
15 each user every day; (b) the vendors request a new secure function evaluation module every day (or a new key), so that
16 vendors cannot use the analysis for more days than they pay. Secure function evaluation computes results from
17 encrypted information without decrypting that information.

18 9. Leveraging Existing Standards

19 The architectural framework outline above can be implemented with a number of existing technical methods. In this
20 section we outline one possible approach, that uses the Extensible Markup Language (XML) to encode web
21 information with meta-tags (that can represent profile information). The Java virtual machine shipped with current
22 Internet browsers can be used to run personalization code, that determines dynamically the information that should be
23 presented to a user.

24 We can use XML to embed the profiling information for products that are offered by vendors, and information that is
25 offered by information providers, directly into the page itself - with semantic labels that allow client-side processing
26 with a Java engine. The Java Virtual machine, implemented on the browser of a user's client machine, takes as input
27 the XML-tagged page of a vendor, and the locally stored profile information that pertains to the user's current
28 pseudonym, and generates personalized content to display on a monitor local to the user.

29 Compare this architecture to a traditional client-server based solution, where the server produces personalized content
30 for the user, content that is pushed to that user. Such a system architecture requires that the server has direct and
31 explicit access to profile information about the user, information that can then be exchanged with other third parties and
32 is out of control of the user. Secure Data Interchange collects and distributes information consistent with user- and
33 vendor-defined policies. Information can be explicit data, including transactional information that is collected by
34 parties that are involved in a transaction — such as the product purchased, and demographic information (e.g. gender,
35 zip code, occupation). Information can also be implicit data, that includes click stream data that logs the information
36 that a user requests and views, and the time-sequence of hyper-links that a user follows as he/she browses across
37 multiple web sites. Profile information is embedded within a web page as metadata, that is data about data - machine
38 readable information that informs an intelligent agent (such as an SDI-enabled browser) about the data that is included
39 in a web document. The extensible Markup Language (XML) proposal of the Worldwide Web Consortium working
40 group on SGML provides an ideal standard for representing such information [XML].

41 XML allows meta-content to be included with documents, machine-readable information that enables documents to be
42 processed by client software. Augmenting web documents with structured information in SDI enables clients to
43 perform user personalization - pushing computation to clients, and allowing greater control over user-profiles because
44 profiles do not need to be released from clients. XML can represent rich data structures, and that allows a grammar to
45 be defined for information that allows data to be automatically verified for correctness [SGML].

46 The ability to embed data within web pages allows client-side processing of information. By embedding profile and
47 location information directly within a web document we can alleviate the bandwidth and computational bottlenecks that
48 can occur at a centralized profile server if profiles are fetched on-the-fly when web pages are downloaded by clients.

CONFIDENTIAL

- 1 The origin server (supported by the vendor) requests periodic profile updates from the central SDI server. This
2 duplication of information enables the profile and the page contents to be provided directly from a vendor's server.
3 There are some potential drawbacks of this approach: (1) the profile information associated with a web page and target
4 objects can be out-of-date; (2) the profile information is available to all clients and proxy servers, not just those that are
5 SDI-enabled; (3) the profile information can be altered. We suggest technical solutions to each of these problems
6 below.

7 9.1 Periodic Update of Web Page Profile Information

- 8 The central SDI server provides profiling information to vendors that subscribe the SDI. A vendor sends a 'Request
9 Profile Update' message to the SDI server, to instruct the SDI server to send new profile information. The SDI server
10 responds with a 'Profile Update' message, that contains updated profile information, generated from explicit and
11 implicit data that it receives from vendors and users in accordance with privacy policies. The request-response
12 mechanism can be implemented using the standard HTTP Post/Response mechanism in conjunction with XML
13 message types. For example, the server 'Request Profile Update' message can be represented in XML as:

```
14 <?XML version = "1.0" ?>  
15 <?xml:namespace ns = "http://www.sdi.com" prefix = "SDI" ?>  
16 <!doc>  
17 <SDI:Request> http://www.some_vendor.com </SDI:Request>
```

- 18 and the SDI server 'Profile Update' message can be represented in XML as:

```
19 <?XML version = "1.0" ?>  
20 <?xml:namespace ns = "http://www.sdi.com" prefix = "SDI" ?>  
21 <!doc>  
22 <SDI:Update>  
23 <SDI:Profile>  
24 <SDI:Item> (1231, 0.453) </SDI:Item>  
25 <SDI:Item> (1041, 0.034) </SDI:Item>  
26 </SDI:Profile>  
27 </SDI:Update>
```

- 28 An illustrative Document Data Type (DTD) for an SDI:Profile element type is presented in the next section. The XML
29 messages are included in the body of standard HTTP Post/Response messages. We limit the performance degradation
30 caused by out-of-date profile information that is stored within web pages of on-line vendors by associating "out-of-
31 date" time stamps with the profiles that are provided by the central SDI server. This mechanism is similar to the
32 "expiration time" tag of a Netscape Cookie message. The frequency with which profile updates need to occur will
33 depend on the speed with which profile information changes. The "out-of-date" time stamp can be included as an
34 additional element in an SDI:Update message.

- 35 Vendors request new profile updates when the current profile information is out-of-date, and more frequently if
36 required (although we allow for a per-update charge). This is a "pull" model for profile-updates. An alternative
37 architecture for profile updates is a "push" model, where SDI periodically sends new profiles to be added to the web
38 pages on a vendor's server. The "pull" model is our preferred model because the responsibility for maintaining current
39 profile information is decentralized, resting with the vendors in the system.

- 40 The system as outlined above can be implemented within the current HyperText Transfer Protocol (HTTP), as a
41 sequence of challenge/response pairs between clients and servers. The HTTP Post/Response mechanism allows clients

CONFIDENTIAL

and servers to exchange data, and this data can be an instance of an XML Document Type, within the body of a HTTP message. The HTTP protocol is the underlying mechanism, with SDI messages contained in the body of the HTTP Post and HTTP Response as XML documents. In one variation of SDI the profile of a user is maintained on the user's client, and partitioned into separate profiles for each pseudonym that a user chooses to maintain. Personalization of products and services (product types, prices, etc.) is performed at the client, through the execution of trusted code that is embedded as a Java applet or as JavaScript within the web document of a vendor. In this way a vendor never receives access to the profile of a user, but is nevertheless able to personalize its response to users, even when a user first visits a site (on the basis of the profile for a user from his/her previous online transactions). Profiles for the target objects of a vendor that enable appropriate objects (representing particular products, or news stories for example) to be presented to a user are embedded as XML data within the vendor's web document.

In another variation of SDI personalization is not performed at the client, but either at the ISP-level SDI proxy server or the vendor's server. The location and other profile information that relates to a user are pushed to the ISP-level proxy or vendor server when a user requests a web page. In the same way as XML allows profile information about web sites and vendor products to be associated with a web document, and profile information to be provided from the central SDI server to a vendor, XML can be used to encode a user's profile. The system of SDI allows for profile and location information to be randomized slightly (and even anonymized) to protect the identity of a user, for example when an ISP-level proxy is not trusted.

9.2 Example: A Possible XML Representation of a User Profile

The World Wide Web Consortium (W3C) SGML working group developed XML (extensible markup language) to provide an open and extensible grammar for structured data [XML]. An XML document has an associated schema definition to enable an XML-enabled browser to validate the structure of XML data automatically. A Schema in XML is called a Document Type Definition (DTD), and defines the names of tags, their structure, and their content model. XML allows the DTD for an XML file to be identified through a Universal Resource Indicator [URI] in the header of the file (see below). XML also allows URIs for mobile code resources to be referenced, in order to enable a client to process embedded XML data. An XML document must be well formed, and in order to be well formed the tags must form a tree structure. In addition, the DTD allows the structure of an XML document (an instance) to be validated against a particular schema. Senders and receivers must only send valid SDI files. Each SDI message is a valid XML document.

We provide an example XML instance and part of a Document Type Definition for use within the system of SDI. Profile information, as generated automatically through collaborative filtering techniques (for example, see issued US Patent #5,754,939) can be represented as a list of attribute-value pairs within an XML document. An attribute is defined by a numeric code, and the value defines the weight of the attribute. For example:

```
<?XML version = "1.0" ?>
<?xml:namespace ns = "http://www.w3.org/OPS/OPS" prefix = "OPS" ?>
<?xml:namespace ns = "http://www.sdi.com" prefix = "SDI" ?>
<!doc>
<SDI:ProfileData>
  <SDI:Location>
    <SDI:Geocode> 12321561 </SDI:Geocode>
    <SDI:DigiMap> http://www.digimap/?12321561 </SDI:DigiMap>
    <OPS:Zip> 19103 <SDI:/Zip>
  <SDI:/Location>
  <OPS:Demographic>
    <OPS:Gender> F </OPS:Gender>
    <OPS:Age> 26 </OPS:Age>
    <OPS:Income> 50000-75000 </OPS:Income>
  </OPS:Demographic>
  <SDI:ID>
    <SDI:Pseudonym> P12543 </SDI:Pseudonym>
```

CONFIDENTIAL

```
1      <SDI:PublicKey> 12453246129421 </SDI:PublicKey>
2      </SDI:ID>
3      <SDI:Profile>
4          <SDI:Profile-item> (1242, 0.546) </SDI:Profile-item>
5          <SDI:Profile-item> (56, 0.045) </SDI:Profile-item>
6      </SDI:Profile>
7  </SDI:ProfileData>
```

8 The Document Type Definitions for this document are specified in the header, and include URIs to a DTD of the Open
9 Profiling Proposal of the W3C, and also a DTD of the Secure Data Interchange. The OPS DTD is used to boot strap the
10 SDI DTD, providing tags for common profile information, such as 'Gender', 'Age', 'Income', etc. The section of the
11 SDI Document Type Definition that is used in the above XML fragment is presented below. It makes reference to tags
12 defined in the OPS DTD, and the RDF (Resource Description Framework), a W3C proposal to standardize the structure
13 of DTDs for XML documents. XML Name spaces [NS] provide a method for unambiguously identifying the
14 semantics and conventions governing the particular use of property-types by uniquely identifying the governing
15 authority of the vocabulary, for example OPS and SDI in the example above. The URI for a schema can contain a
16 human and machine-readable description of an XML schema.

```
17      <!ELEMENT SDI:ProfileData (SDI:Location?, OPS:Demographic?, SDI:ID?,
18      SDI:Profile?) >
19      <!ELEMENT SDI:Location (SDI:Geocode?, SDI:DigiMap, OPS:Zip?, OPS:Address?) >
20      <!ELEMENT SDI:ID (OPS:Name?, SDI:PublicKey?, SDI:Pseudonym?) >
21      <!ELEMENT Profile RDF:list<SDI:Profile-item> >
22      <!ELEMENT SDI:Geocode #PCDATA >
23      <!ELEMENT SDI:Digimap #URI >
24      <!ELEMENT SDI:PubicKey #PCDATA >
25      <!ELEMENT SDI:Pseudonym #PCDATA >
26      <!ELEMENT SDI:Profile-item (SDI:Attribute-ID, SDI:Attribute-value) >
27      <!ELEMENT SDI:Attribute-ID #PCDATA >
28      <!ELEMENT SDI:Attribute-value #PCDATA >
```

29 The tag '#PCDATA' is used here to represent numeric or textual information, '#URI' declares that an instance of
30 element 'SDI:Digimap' must be a valid URI pointer.

31 9.3 Maintaining the Integrity and Security of Embedded Profile Information

32 The privacy of information in transit between servers and clients can be assured through standard end-to-end
33 cryptographic solutions that establish a secure session prior to any data exchange, such as Secure Sockets Layer (SSL)
34 that uses X.509 certificates and is supported by current browser technology.

35 In order to prevent the possibility of individual users being bribed by vendors to disclose target object profile data
36 which reflects this type of information, users should not be provided access to directly decrypt the metatags for these
37 portions of the target objects profile data, but rather this decryption and release of profile data should be performed
38 securely in conjunction with the functions of the profile processing (profile matching module) upon the client level
39 proxy server rather this decryption and release of profile.

40 In addition, we prevent unauthorized access of embedded profile information through the encryption of the metadata
41 that is represented within the XML structure of a web page. Profile information can be encrypted using a hierarchy of
42 keys, so that different levels of access to the information may be provided according to the access levels of users and
43 vendors. All users that request web pages from SDI-enabled vendors, whether or not the user is a member of SDI
44 receive the same profile information. We provide encrypted profiles to vendors in the 'Profile Update' messages from
45 SDI to vendor servers, so that: (a) unauthorized agents cannot tamper with the profiles; (b) the profiles cannot be read
46 by unauthorized agents.

47 The SDI system supplies a private key to trusted SDI client software, that enables only SDI-enabled clients to access
48 profile information, and only access that information to the extent permitted by privacy policies of users and vendors.
49 Different levels of encryption enforce multiple levels of access. Periodically the key pairs are changed to prevent

CONFIDENTIAL

extended attempts at cryptographic attacks. The SDI system uploads the key that provides the correct level of access for a user to a user's client, once terms of access and profile management have been agreed. A client can only access embedded information once enabled with a relevant key. Finally, profile information is signed with a digital certificate, to prevent third parties from tampering with profiles for commercial gain.

9.4 Related World Wide Web Consortium Proposals

The proposal for a Meta Content Framework (MCF) suggests a particular structure for the description language for web pages, to enable schema to be shared and re-used [MCF]. This proposal is incorporated into the W3C Resource Description Format standard [RDF]. The proposal for an Open Profiling Standard [OPS] describes a system for profile exchange between two parties, building on XML and MIME standards. The W3C SGML working group has defined XML to provide an open, extensible grammar for structured data. The proposal on privacy and profiling on the Web [PRIVACY] extends the vCard [VC] schema for electronic business cards to include profile information, and suggests that profile information can be stored and managed locally, with Client-server exchange of personal information as required (using the HTTP challenge/response mechanism).

The Resource Description Framework (RDF) enables the encoding, exchange, and reuse of structured metadata. RDF is an application of XML, with additional constraints to allow for DTDs to be published, and interchangeability across different communities. The ability to standardize the declaration of vocabularies will encourage the reuse and extension of semantics among different information communities [Mil98]. RDF is a W3C proposed standard for defining the architecture necessary for supporting web metadata. RDF is an application of XML that imposes structural constraints to provide unambiguous methods of expressing semantics for the consistent encoding, exchange, and machine processing of metadata. RDF additionally provides means for publishing both a human-readable and a machine-processible vocabularies designed to encourage the exchange, use and extension of metadata semantics among disparate information communities.

10. Support for E-Commerce Functionality

10.1 Generation of Mailing Lists

We can use the same profile information that provides focused/personalized service to users that hit a site that they have not visited before to form well-targeted mailing lists for vendors. The Secure Data Interchange can form mailing lists in a number of different ways.

First, consider a vendor that wishes to send targeted mail to some of its own user-base. When users connect to a site they indicate whether or not they are willing to receive electronic mail, and provide a "mail certificate" to a vendor if they are happy to receive mail. The Secure Data Interchange can proceed as follows:

- (a) perform analysis for the vendor to determine an appropriate set of users to receive the solicitation, based on the information that the vendor provides about what it intends to market, and provide the list of pseudonyms to the vendor for mailing; (b) perform the same analysis, but also forward the communication to the users directly.

Now, consider a vendor that wishes to target new users, represented with different pseudonyms. Users indicate whether the information that a vendor submits about his/her transactions may be used for solicitations, and furthermore vendors indicate the set of business interests that can receive the benefit of information that is submitted to the central SDI server. The SDI server can continue by performing analysis on the relevant subset of the permitted class of data records that pertains to the product or service that the vendor wishes to model, and generate a list of appropriate pseudonyms. Finally, the SDI server can sell the pseudonyms to the vendor outright, together with a certificate that the vendor can send mail to the pseudonyms, or the SDI server can retain control by sending the mail on behalf of the vendor.

Provide vendors with virtual mailing lists that can be mailed to via the proxy server only, i.e. these customers should be solicited based on our analysis. (could even give summary info., without revealing details about users). Furthermore, do not even reveal data that corresponds to a pseudonym to a vendor because the vendor then has that information about me when I am on his site.

CONFIDENTIAL

A central data warehouse also enables vendors to identify new potential customers. This process is broken down into a number of steps:

- (i) The vendor assesses the value of the information present in the secure data interchange. This computation is performed securely either by revealing randomized aggregates to the vendor to enable its own local analysis, or by allowing the vendor to check data and algorithms into the secure data interchange site for analysis.
- (ii) The vendor selects criteria for mailing unsolicited advertisements, and agrees on a pricing model. In this case per-impression pricing is the most obvious pricing model, as it is difficult to monitor when a user responds to unsolicited mail per-transaction pricing is difficult. The user could be motivated to do this should the Secure Data Interchange promise future returns for recording a successful solicitation with the database.
- (iii) Either the data list is released to the vendor for its use, if this is within the selling vendor's data policy, or the data interchange sends mailings on behalf of the purchasing vendor.

10.2 Physical Mail

(a) Vendor to User

Figure 13 shows a flow chart for the process of sending physical mail from a vendor to a pseudonymous user.

A vendor must hold a "physical mail certificate" to be able to send mail (packages, letters) to a user under a pseudonym. The certificate is similar to the "electronic mail certificate", in that it is signed by the private key of the user's pseudonym, and indicates that the vendor with public key P^*V can send mail to the user (under the pseudonym).

Each user has a trusted physical address authority, just as it has a trusted electronic mail authority (the second-level proxy server), that maintains the physical mailing address for each pseudonym. When a vendor has a letter X to mail to user with public key PKP , the vendor generates a unique ID for the package, IDX , and sends the ID code and the physical mail certificate to the trusted physical address authority of the user.

The physical address authority receives the certificate, $S(PKP, PK^*V, SEND-MAIL)$, SKP , that indicates that the vendor is authorized to send mail to the pseudonym, and the packages identify code, signed by the vendor to certify that the vendor holds the secret key that matches the public key in the physical mail certificate.

The vendor then passes the letter X and the signed ID code to a trusted mailer, that supports pseudonymous mailing, and has been certified by the central SDI server as such. The trusted mailer then provides the signed ID code to the physical address authority, signed with the private key of the trusted mailer. The physical address authority verifies that the trusted mailer is a valid service, and releases the real address of the user to the mailer. The mailer now has the letter X that the vendor wants to send to the user with pseudonym P , and the physical mailing address of the user - and the package can be mailed. At no time did the vendor determine the true mailing address of the user, unless it works in collusion with the trusted mailer, but the trusted mailer is certified by SDI, and also audited by the chosen physical address authority of the user. The address authority will only release addresses to reputable pseudonymous physical mail agents.

We can operate physical mailing lists in the same way, and gain additional security by never releasing the pseudonyms or the mailing addresses to the vendor that has requested the targeted solicitations. We can use a technique that is similar to the technique that we used for virtual mailing lists. The vendor describes its solicitation to the central Secure Data Interchange, which leverages as much data as possible (without violating the privacy policies of any of the users or vendors that are represented within the data). The central SDI server generates a list of suitable pseudonyms, and then provides a series of unique codes to the vendor, that the vendor can supply to its chosen pseudonymous mailer with the material that is to be mailed. The central SDI server also provides the appropriate address authorities with authorization to release the physical mail addresses to the mailer when presented with the IDs. Notice that at no stage did the vendor have the pseudonyms or the mailing addresses. The parties all have only as much information as is necessary - the vendor needs somehow to identify its packages to the pseudonymous mailer. The mailer needs an identifier to present to the address authority, and receives the addresses. The address authority just needs to know what addresses to release and to which third party.

CONFIDENTIAL

(b) User to Vendor mail

The Secure Data Interchange system also provides a mechanism for users to send physical mail to vendors that are registered with SDI with pseudonymous return addresses. In particular, when a user sends mail to a vendor, the first-level proxy server provides a tool that: (1) Computes/Looks-up the appropriate pseudonym for the user with this vendor. (2) Generates a unique ID, and submits a signed message to the central SDI-server, where the message relates the pseudonym, the vendor, and the ID. (3) Provides the unique ID to the user.

The user writes the unique ID on the envelope, and mails it to the vendor. Should the vendor wish to reply to the user, then the vendor can take the envelope to a pseudonymous mailer, and request that the envelope be mailed appropriately. The pseudonymous mailer verifies the identity of the vendor, and then submits the ID, together with the vendor's signature, and its own signature, to the physical address authority that is maintained by SDI. SDI releases the address to the mailer that can then return the mail.

10.3 Pseudonymous Payment Mechanisms

The Secure Data Interchange architecture must be able to support all the standard electronic commerce functions that we take for granted, but while maintaining pseudonymity for users and following privacy policies. There are various different solutions to this problem.

(a) Anonymous Credit Card Payment

The second-level proxy server can maintain information on the user's credit card information, and perform the following transaction. Whenever a user makes a purchase from a vendor, the user provides the vendor with authorization to bill \$x to his/her credit card account, but anonymously - through the Secure Data Interchange as a middleman. The user generates a unique number, Y, and signs a "right to payment" message, $M = (\$x, PKP, PKV, Y)$, that gives the vendor the right to make a claim for payment of \$x from the Secure Data Interchange. The first-level proxy server registers the unique number Y with the second-level proxy server to ensure that the vendor does not spend the money twice, and provides the proxy server with authorization to charge \$x to his/her credit card when the request for payment is presented.

When the vendor submits its "right to payment" and proof of identity to the second-level proxy server the proxy server first runs the charge through the user's credit card, and if that clears, runs the charge from the vendor through the account of SDI (which could also be a credit card, or could be operated as electronic cash or some other mechanism for payment).

This "anonymous credit card" payment method has the following properties:

1. The user's credit card pays \$x, but does not know who receives the money except that it is going to the Secure Data Interchange.
2. The vendor receives payment for \$x, but does not know the user's credit card information, or the user's identity.
3. The Secure Data Interchange incurs no financial risk because it receives payment from the user before making payment to the vendor, although there could still be problems if the user complains about the quality of the good for example.

This protocol is simpler than full cryptographic anonymous credit card mechanisms because the SDI acts as a trusted third party to both the user and the vendor. << provide literature references >>

(b) Electronic Cash

Electronic cash is anonymous, just like physical cash. The user purchases electronic cash from an electronic bank, presenting blinded notes, so that the bank has no record of the note numbers that it issues to the user. For example, the user generates a new note number, X, and has the bank sign a blinded copy with its \$10 signature, $S(B(X), SK_{BANK\$10})$. Then the user, or the first-level proxy for the user, removes the blinding factor, and can use the

CONFIDENTIAL

1 electronic cash as tender. Whenever the note changes hands the recipient needs to check with the bank that it has
2 not yet been spent, because notes are easily copied, but not forged.

3 Electronic cash has the same useful properties as anonymous credit cards, although it is perhaps a little more
4 exotic. In particular, notice that the bank does not know to whom, or for what, payment has been made. and the
5 vendor does not know which user made the payment - it just receives the payment. We have minimized the amount
6 of information exchange that takes place between the various parties in the system.

7 11. Models of Data Release

8 Data can be released with associated "terms of disclosure", that define: (a) the price or "exchange rate" of data; and (b)
9 the time and usage parameters of access to the data. Possible time and usage parameters include: data may be used
10 indefinitely, a periodically renewable usage contract, a number of impressions contract, a number of transactions
11 contract, as a one-time permanent or temporary exchange, as a time limited privilege, or subject to certain conditions.

12 Data can be priced according to its accuracy and content. For example, aggregate data might be sold more cheaply than
13 detailed data, and we can sell the right to access data dynamically (on demand), on a per-impression basis, or up-front.
14 Some vendors might wish to purchase the rights to utilize the data. The vendor might also be interested in renting data
15 (for example, when the value of data is uncertain), and purchasing the data outright if it proves valuable.

16 The central data warehouse of the Secure Data Interchange can operate as a competitive market place for data and
17 information, with the possibility of facilitating the trade of information between vendors that have synergies in their
18 data sets, perhaps with SDI acting as a valuation device to provide information about the likely benefits of an exchange.
19 The form and terms of data exchange will be those that provide greatest benefit to the contributing vendors.

20 In general, vendors will most likely be interested in purchasing the access to profile information about users that hit
21 their web pages dynamically, on a per-impression or even per-transaction basis. This removes much of the uncertainty
22 from the vendor about the coverage and relevance of data that is provided.

23 The methodology that we promote within the basic SDI architecture, where users themselves maintain control over
24 their profiles also provides users with ownership of that information, and enables users to maintain ownership while
25 providing per-impression access to vendors, who receive the results of using the profile to personalize an interaction
26 but not the profile itself. The profile need never be released by the user.

27 Per-transaction pricing, possible in on-line applications of secure data interchange, offers powerful new modes of
28 business. For example, it is possible to monitor the number of "click-throughs" achieved on a particular banner ad,
29 and charge the requesting vendor on a per-transaction basis. This provides a self-enforcing structure to a contract. It is
30 in the best interest of a providing vendor, and the secure data interchange, to provide accurate data to enable proper
31 focusing of ads, and also to provide good data analysis tools, because the success of the advertising campaign
32 determines the revenue that they receive. It is as though the providing vendor is working on an "on-commission" basis.

33 Ideally we would like to sell data to a vendor at the value of that data to the vendor. This should be done in a way that
34 prevents the vendor from selling the data on to another vendor on the black market. One solution to this problem is to
35 "rent" the data to a vendor, but not give a purchasing vendor outright control of the data. The secure data interchange
36 can be used to provide access to the data. For example, instead of physically transferring data to the data warehouse of
37 a purchasing vendor one could provide a mail-distribution service for that vendor at the data interchange. The vendor
38 could specify constraints on the users that are to receive a particular solicitation, give the interchange "Mailer" the
39 advert, and then request that the interchange deliver the solicitation.

40 The interchange could also provide a vendor with the right to run analysis at the interchange, without actually down-
41 loading the data.

42 12. Variations on the Basic SDI Architecture

43 There may be commercial contexts in which an SDI service can be established where there is already in place a pre-
44 existing trust relationship between multiple vendors and a third party. Such third parties are inherently motivated to
45 provide services to enhance advertising and e-commerce for their existing and potential customers. These third parties
46 may include, for example, web hosts or e-commerce service providers (ESPS) which often have hundreds or thousands

CONFIDENTIAL

1 of sites which they host, Web portals, information and commerce service manufacturers, advertising and affiliate
2 network services and data analysis and business intelligence tool providers (which includes the business to business
3 application).

4 A third party may wish to implement an SDI which operates separately and independently from the central SDI
5 service. Alternatively, some of these third parties may install an SDI server on their customer information server. The
6 server may be integrated into an existing advertising service which they operate and maintain, in which case the vendor
7 receives an appropriate fee for data which is exchanged between his/her existing customers, and a reduced fee (which
8 may be split with the central SDI service) for data which is exchanged by/between a member of his/her SDI service and
9 vendors who are members of the central SDI service but not of his/her local SDI service.

10 12.1 An open SDI system

11 An ISP level proxy server can contain the user profile generation module, profile processing module, user profile
12 interest summary generation module and target object generation module which operate in distributed manner. This
13 enables an ISP to independently implement the core functionality of the system without the cooperation of information
14 vendors (Web sites) or their operators (Web hosts). The modules in third-party SDI servers can share information with
15 the modules in network vendor servers. This flexible architecture enables ISPs to implement SDI and when available
16 also the complete data sets available from the information vendors.

17 SDI can allow third parties to operate their own secure advertising and/or electronic commerce-based product
18 syndication affiliate network (for all customers). The users that are also subscribed to SDI, can be given highly
19 personalized information for each site or for the network of hosted sites (which could involve an interface which
20 provides site to site links as a "virtual mall"), and a menu interface to these sites which includes the 2 or 3 dimensional
21 personalized menuing features and personalized search facilities as disclosed in the parent description (a "personalized
22 portal").

23 12.2 A Closed SDI System for a Syndicated Network

24 The Web host (or more generally a vendor, a provider and/or operator of server functionality to a variety of information
25 vendors), may also be interested in operating his/her own closed version of SDI. The main SDI server for the closed
26 system can be located on the network vendor servers, or it may reside upon the information vendors servers (as it is
27 operated by that local Web host). For example, an affiliate ad network (including a web host acting in such capacity)
28 could upon installing SDI onto their network enable and enforce the wishes and desires of advertisers (and particularly)
29 sites which are advertised upon with regards to what types of sites and advertisers (respectively) they allow or disallow
30 for purposes of standard or affiliate advertising, in accordance with the methods herein disclosed. The general
31 implementation for determining which this general application for using collective user feedback to determine relevant
32 site links was described in the parent issued patent). In this case, end users who are subscribed to SDI would receive
33 personalized affiliate links (including product level recommendations for on-site purchases) which have been pooled
34 and profiled at the main SDI server from all SDI vendors (in distributed fashion) and matched with the user.

35 12.3 Interoperability Between Local SDI Services

36 With interoperable (local) SDI services, we can also facilitate the secure enforcement of data sharing policies and
37 transfer of transaction fees between these local SDI services E.g., by/between aggregations of ad networks, syndication
38 networks and Web hosts operating virtual portals and advertising/syndication networks with personalization as its
39 primary capability.

40 In each of these primary example domains, the server operator is financially motivated to sell the SDI services to
41 his/her sites because the transaction based model is used, and the server operator receives the commission on each
42 transaction (or click through) occurring within his/her network of sites. However, if the server operator also integrates
43 his/her local SDI service into the main SDI service (to share user lists and impressions and/or space to advertise to
44 these target users), s/he can receive a commission (in conjunction with each vendor transacted with) for each
45 advertisement placement or syndicated transaction to or from his/her network.

CONFIDENTIAL

- 1 We can also allow the local server operator to split the transaction fee (normally received from the main SDI service),
2 thus "referral fee" for both the referred customer and the referral of customers (through the placement of outside ads or
3 products on one of his/her sites) or other means of targeting his/her site's existing customers.
- 4 Reduced overhead resulting from economies of scale which may likely result in incentives to the local operation, e.g.,
5 free installation and operation of his/her local main SDI server by the main SDI service, i.e., as the operational
6 overhead would be cost justified by the shared transaction fees of customer referrals and advertising space coming back
7 to the main SDI service.
- 8 This architecture also may be useful and is ideally suited for cross vendor product advertising as through an ad network
9 or product syndication network using affiliate links. In addition to the user profile generation module, a target object
10 profile generation module should also reside across the network vendor servers such that it is possible to generate target
11 object profiles for target objects on network vendor servers. Alternatively, user profiles and target object profiles are
12 downloaded to the client level proxy which performs collaborative filtering tasks as the user browses from site to site.
- 13 In both of these cases, the main SDI server can receive user profile data generated from the user profile generation
14 module located on the ISP-level proxy, and target object profiles generated from the target-object profile generation
15 modules located on the various multiple information vendor servers.

16 13. *Dynamic/Real-time Secure Data Interchange*

- 17 The user-centric SDI model allows users to provide personal information on a carefully controlled basis to vendors and
18 other users. Furthermore, vendors can implement rules that personalize the information, products, and services
19 provided to users—on the basis of personal information that they receive from users directly, or have acquired about
20 users.
- 21 As an extension to this model, we also allow users, vendors, and other third parties to associate "meta-information"
22 with other users and vendors. This information might be a user's opinion about his/her interaction with another user, an
23 annotation that relates to a particular web page, or information about a physical object, for example how to get to the
24 top of a tower. The system of SDI enhances the value of this information by providing a secure environment where
25 users can also associate their own profile with the meta-information that they "leave". This allows collaborative
26 filtering techniques to generate appropriate meta-information about an object (user, physical object, vendor, web page,
27 etc.) that will be useful to a particular user—given that user's own profile.
- 28 We define "virtual tags" as any piece of information about an object (physical or virtual). The information may be
29 authored by any party, but annotated accordingly. For example, the appropriate virtual tag provided by a user about
30 his/her-self is the pseudonymous profile for that user, — and with SDI only the user his/her-self can gain access to the
31 profile (either directly through editing, or indirectly through continuing transactions).
- 32 It is useful to implement a "reputation system" within such a virtual community. Initially users (under pseudonyms)
33 have no reputation, and their opinion does not count for much, but after every positive interaction (as defined by other
34 parties in an interaction), the "reputation" of a user can increase. (see the Kasbah system, MIT). This reputation system
35 is appropriate to a pseudonymous environment. Notice that gaining negative reputations is not useful when users can
36 simply change identities.
- 37 In one variation we can "block" certain users from providing information, when those users have negative reputations.
38 Clearly, collaborative filtering or other data mining techniques could usefully allow for reputations when weighting
39 information about an object.

40 13.1 *Autonomous Exchange of Information*

- 41 Client-level SDI proxies can act as autonomous agents in an architectural variation of SDI, where the "client-level
42 proxy" is co-located with a (physically) mobile user, for example on a palm-held computer or head-up display.
- 43 In a "match-making" application domain, the goal of user-agents is to find other user-agents with desired
44 profiles/synergies, and arrange person-person meetings, or business-business meetings/agreements. There are two

CONFIDENTIAL

1 parties in any exchange of information, and although information exchange will be bidirectional, it is useful to talk
2 about a "requestor" and a "requestee". The key provision provided within SDI is that:

- 3 1. Requestors and requestees can communicate anonymously, without revealing (even pseudonymous) identities.
- 4 2. A requestor cannot access profile information about a requestee unless authorized to do so (implicitly or explicitly)
- 5 by the requestee.

6 Implicit authorization occurs when a requestor can present certificates to verify that it has required attributes to access
7 particular information. Explicit authorization occurs when a requestee provides direct authorization to a particular part
8 of a user's pseudonym.

9 Essentially there is bidirectional information filtering: the requestor agent will only present certain information to the
10 user, information that is relevant; and the requestee will only provide information when a request is judged to be
11 legitimate. Information exchange between agents occurs as part of a multi-step negotiation, until both parties can agree
12 on terms for either a physical meeting (or execution of a deal), or further pseudonymous exchange of information or
13 cooperation.

14 In this case, using methods taught in the co-pending patent << LEIA >> user-agents can identify other agents that are
15 "close" through an anonymous matching market, where agents provide their location and a (one-time) identity for
16 contact. The market informs agents when other agents are close. In fact, this "anonymous matching market" can be
17 extended, and agents can provide more than just location information.

18 A user-agent might also broadcast a "persistent" query over the agent-network, for example requesting response from
19 agents with a particular set of attributes, and providing some information. Decisions about what information to
20 exchange are made on the basis of both static and dynamic profile attributes, e.g. standard (historic) profile
21 information, current behavior, current location, recent activity, and credentials that can be presented/denied. LEIA
22 style-behavior attributes can be used to automatically decide on the relevance of new virtual tag information. A
23 requestee might also demand certain credentials to indicate the lack of negative reputation marks, for example that an
24 interaction with the user has never received a bad rating. Perhaps a third-party could be used to determine whether the
25 user's know each other (eg www.sixdegrees.com)

26 We can extend information disclosure to include communications between users and other parties. A user-agent might
27 decide to make another agent privy to communication, on the basis of the context of communication, content, location
28 of parties, profile of the third-party, etc.

29 When a requestee denies a request for information, it may instead provide criteria for data releases. A requestor can
30 respond with a different information request, or a subset of required credentials. Finally, the agents might agree on
31 terms of negotiation and conditions can be anonymously fixed.

32 Negotiation might FAIL, in the case of missing rules in a rule-set, or a special case that has not been anticipated. In this
33 case direct user-user interaction might be necessary, although this can be executed anonymously via a real-time
34 anonymizing service.

35 There are (at least 5 levels) of information disclosure:

- 36 1. indicate to another user interest
- 37 2. release of profile information
- 38 3. disclose communication
- 39 4. add an individual to a current correspondence session
- 40 5. schedule a meeting/strike a deal

41 The end-result of information exchange could be an agreement to calendar a meeting for some future time and place;
42 and absolute, or pseudonymous revelation of identity.

43 Initially, an implementation of the data-release policies might allow only manual definitions. However, after an initial
44 "beta testing" phase, a data mining suite could be used to cluster users and generate exemplar data release and data
45 request policies. A system can provide default settings for users, and recommend setting based on users with similar
46 profiles. The user can further fine-tune the rules.

47 Finally, automatic feedback techniques can be useful to adjust rules, for example—when a user is especially receptive
48 to particular type of introduction then make such introductions more likely in the future. An intelligent interface system

CONFIDENTIAL

- 1 might also suggest refinements to the rules, to automatically cover "patches" where the user currently controls
2 interactions. (see the Microsoft Lumiere project).
- 3 Note 1: must next describe user profile attributes and rating criteria, e.g. selective access to group ratings and
4 annotations by navigating their attributes. Also automatic generation and comparison ratings between items via rapid
5 profiling.
- 6 Note2: Real time (including spoken) communications (either user to N-users or user to user) within a local proximity
7 where the users are "pre-qualified" (e.g. knowledgeable in that particular field, educated, etc.)

8 13.15 Credentials: Examples and Typical Uses

- 9 In accordance with the parent issued patent, various credential issuers are provided for issuing standard and resolution
10 credentials to individuals. Thus certain entities may be entrusted with "legitimate authority" to validate and submit
11 credentials which are issued to the appropriate individuals. Resolution credentials are provided to prove the absence of
12 a quality attribute or behavior (which is often of a negative nature) relating to an individual and is submitted by a third
13 party and typically must be issued on a periodic basis in order to maintain currency. If a resolution credential is not
14 issued (or not renewed) an adjudicating third party is provided which has access rights to both of the parties is provided
15 to resolve resulting disputes (from the subject user). The present invention describes how credentials can be issued to
16 users pseudonymously.

- 17 A few simple examples of resolution credentials which may be of interest to users (credentials which users may
18 commonly request as a precondition to requesting or accepting requests to be introduced or initiate communication with
19 an outside unknown third party) include:

- 20 1) (for business associations) are in good business standing, e.g., have not attempted to defraud other users in the
21 course of common business practices. Or maintain sufficient funds in one's account to perform business activities
22 (as represented by the user).
- 23 2) (for business interactions or social interactions) are in good standing with the law.
- 24 3) are considered "safe" individuals, e.g., have not been rightly accused or criminally convicted of violent acts, (e.g. a
25 store clerk may wish to be made aware when an individual walks into a store which cannot present a resolution
26 credential for not having been convicted of retail theft or robbery.
- 27 4) have not been accused by other individuals of inappropriate or antisocial behavior.

- 28 Some standard credentials which may be of interest to many users, and which may (as with resolution credentials) be
29 incorporated with the standard settings of the user's data request policy as herein described. A few examples are cited
30 (among countless potential others): profession, awards, honors, alma mater, e.g., Harvard graduate, doctorate degree,
31 etc.

- 32 There are a variety of rules which a user's data disclosure policy and data request policy may contain, to control what if
33 any attributes are released, and what credentials are required. A data request policy may state a rule for explicitly
34 notifying the user if a particular resolution credential (e.g., indicative of a serious problem or concern)

- 35 cannot be presented in response to the
36 user's disclosure request.

- 37 Some users may not wish to disclose specific information about themselves via these standard credentials but instead
38 certain "extracted" more general information may be provided about themselves. For example, instead of a "Harvard
39 grad or Ph.D." there may be, for example, credentials indicating "intellectual" or "prominent intellectual". Or instead
40 of indicating an individual's wealth or value of assets, the credential may indicate "wealthy" or "very wealthy"
41 (typically, depending upon user's wishes this latter credential should also be withheld during initial introductions or
42 subject to some fairly stringent conditional criteria from the other party) and instead replaced with an even more
43 general credential e.g., "prominent" or "influential citizen"). Similarly, an individual's exact profession or scope of
44 work may not be fully disclosed initially but rather a more general definition of his/her profession or perhaps the
45 general field initially in which the user works.

- 46 Another example of a credential of potential interest may include the profiles of users which a certain individual
47 associates with or is acquainted with. The ability of a third party to gain access to this information, however, is

CONFIDENTIAL

1 conditional upon the data release policy of that associate's or acquaintance's data (e.g. it could be affected by what is
2 the profile of the common acquaintance to whom that user would be disclosed as an associate as well as, importantly,
3 the profile of the prospective discloser.) In one variation, the system may simply identify the fact that there are
4 common associates and acquaintances between the two individuals. Again that associate's or acquaintance's data
5 release policy may further control even detection of this fact. It may instead also notify one of the parties of this fact,
6 but request that it not be disclosed to the other party.

7 In accordance with the parent patent application, rules may be learned regarding certain things that a user does (as in
8 ascribing these rules for which messages to send to whom or what user profiles and under what circumstances/events
9 surrounding the target user). Thus, his/her agent may begin to suggest certain future actions which could be performed
10 in the future upon user approval or even automatically. If the user has had no previous interaction at all with the
11 system, it may identify which other users of the system the present user is most similar, and recommend initial rules.
12 Additional textual attributes can also be leveraged to provide extra criteria, and data mining techniques used to generate
13 more appropriate rules.

14 Another category of user credentials include features that may be inferred implicitly by location/time data captured by
15 LEIA. Such information may reveal a user's likely behavior and activities. These inferences, however, are
16 unavoidably somewhat speculative and inconclusive, thus cannot be substantiated on a valid basis for issuing
17 credentials. The data may be useful in suggesting the present context and circumstances surrounding a user.

18 Additionally, the communications which the user may be presently involved in i.e., the content profile of his/her
19 spoken dialogue and/or other "on line communications" may be used and combined with location/time patterns in
20 order to further infer the circumstances, behavior, and present temporal interest of a user and/or third party for
21 purposes of employing the user's data disclosure and data request policies.

22 ** DP—but how can we do this when users are only pseudonymously identified? How can LEIA map physical user
23 IDS to pseudonymous IDs?

24 There are many potential application contexts in which this present architectural framework could be usefully
25 deployed. For example, it would be extremely useful for users who wish to be re-identified or re-contacted after
26 original introduction -- with unique pseudonyms. A user might wish to withhold all identifying information from
27 his/her acquaintances and associates, except when required for further information exchange. The user can also
28 apply the techniques of randomized aggregates to minimize the chance that one friend or associate will be able to
29 correlate that unique pseudonym with the same individual. A new portion of the use's profile could be combined by
30 users with different information about the user or more importantly, a private piece of information under a pseudonym
31 could potentially be associated with his/her identity. The user agent might use a statistical algorithm to identify a
32 threshold of how much data can be released (including what types of data and to what degree it must be randomized),
33 in order to prevent the user's identity from being compromised.

34 Within a location enhanced context, unless the prescribed range of "proximity" to the user is quite large, securely
35 protecting the user's identity from malicious third party collusion (for purposes of combining unique pseudonyms and/or
36 exchanging data that has been released and entrusted to them) is a harder problem. The system could (most obviously)
37 assume data exchange between the parties will occur and limit the combined disclosure to only that of the most data
38 restricted user in a given location/time domain. The system could alternatively, perhaps "space apart" the number of
39 users within a given location/time context who can access more "restricted" user data (of course the problem goes away
40 if all the discloses have similar disclosure restrictions by that user).

41 We allow initial information exchange to be anonymous, such that information that is released as preconditions for
42 release of further information is not useful. Similarly, so long as initial encounters are anonymous there is no need to
43 withhold information about them from the user.

44 ** DP—this is a cleaner solution than an "ISP-level user agent".

45 Credentials can allow users to identify other users that may pose a threat. This identification may be provided vis-a-vie
46 resolution credentials and/or rating (by third parties). e.g. a user has not engaged in any serious criminal activity,
47 physically harmed another person, or interacted with other individuals who are unable to produce these resolution
48 credentials. Other credentials may specify the nature of an infringement, and its context and severity (e.g. what was the
49 context of a physical assault? Was it performed during a bar brawl, against a friend, a boss, an elderly person, a child, a
50 family member — or at work? In this case, the user agent may, for example, bring to the attention of a prospective
51 employer that the user could not present a credential indicating that they had not previously harmed or threatened a

CONFIDENTIAL

1 former employer. Was it minor or severe? Also, if such individuals (lacking, for example, resolution credential proving
2 the absence of having committed armed robbery) are (or come) within a certain proximity of a user, the user may wish
3 to program his/her user agent to notify the user. The same would, of course, apply to a store clerk regarding customers
4 of this sort or to baggage security personnel at an airport. Or, highway patrollers may be interested (e.g., on certain
5 stretches of highway) in being made aware of vehicles and their locations whose agents are unable to provide a
6 resolution credential proving the absence of a drug conviction.

7 ** DP—in general we cannot expect other users to be "broadcasting" negative credentials. The most that we can
8 expect is that a user with negative credentials can not provide (false) positive credentials, and assume that the lack of
9 positive credentials implies the presence of a negative credential.

10 ** DP—think about infringements in one pseudonym can prevent a user acquiring a positive credential under another
11 pseudonym

12 In a more benign though similar example application to that above, users who are "common individuals" could use the
13 above techniques (through similar co-operation/collusion) to "single out" other individuals who may be of a "high
14 profile" nature (e.g., famous or wealthy) who may otherwise strongly prefer to remain incognito, wherein potential user
15 to user communications could result in a significant invasion of privacy.

16 In another application (in accordance with the auto insurance risk determination methods described in co-pending
17 patent application entitled "Applications for Location Enhanced Information Architecture"), an on-board computing
18 device within a user's automobile could identify another automobile lacking, for example, a resolution credential for
19 safe driving, i.e. the on-board user agent continuously polls agents in other cars for a "safe driving" credential, and if it
20 fails to receive such a credential it issues a warning to the user. As an extension, this location data could be converted
21 into a dynamic 2-D rendering upon the user's windshield (using heads up display technology) in order to thus
22 superimpose a persistent flagging or highlighting of that particular automobile from the driver's visual perspective.
23 Pedestrians could also receive instant notification.

24 We can avoid user-identification/privacy violations within such a physical system by fuzzing the information provided
25 to a user, for example indicating only a general location, or hiding the individual within other individuals. When highly
26 sensitive information is disclosed, it is important that we can protect the real identity of the user.

27 Some examples may include: health problems, like HIV status, neuroses, sexual abuse as a child, family history of
28 violence, alcoholism or emotional disturbance. However, a prospective partner might wish to receive credentials of a
29 lack of these (negative) qualities before pursuing (non-anonymous) contact. As is described in LEIA, a roaming
30 cellular connection, or GPS, is not essential for providing a user identifier. For example, optically-based biometric
31 identification techniques such as iris scanning or combined iris/facial identification techniques may be used among
32 other potential inputs as well. Users will be reluctant to release location/time data, even anonymously, when
33 suspicious behavior can be inferred—probably subjectively. Of course, we allow the user to control this data release
34 within his/her privacy policy.

35 Indeed, if a "significant" threshold of suspicious behavior is exhibited and detected, the information may be accessed by
36 law enforcement officials, through seizure of the decryption key for that data (which includes his/her physical location
37 information) and any additional profile data which is considered of immediate critical relevance to the suspect (or
38 prospective) infraction. Such cryptographic techniques for key seizure from a key escrow are well covered in the
39 literature. If behavior is inferred from LEIA information, then the threshold should allow for statistical errors. There
40 may also be certain circumstances in which key seizure may be required after the fact (at some time in the future).

41 For example, if/when certain even moderately "suspicious" behavior patterns are detected, it may be possible for the
42 SDI data warehouse to preserve a comprehensive record of that information (and perhaps the record of that user which
43 precedes and follows that period of interest). Thus preserving evidence which may later prove useful in contributing
44 evidence towards a conviction, acquittal, e.g., proving that a user was not at a particular location/time. A record
45 containing more detailed segments of a user with a proven negative or questionable history may be preserved and
46 general location/time features may be abstracted for the remaining portions of the record (thus compressing the record
47 substantially). This may be performed for regular individuals as well, thus retaining key relevant features while
48 discarding the majority of the record which is irrelevant or redundant.

CONFIDENTIAL

1 A very important caveat underlying the above architectural framework is that until such time as users become
2 personally "wired" (through miniature computing devices and/or wearable computing devices), the above applications
3 involving the use of resolution credentials within the context of a location enhanced (physical) environment will be
4 hard to implement practically. Automobiles may be an exception, as may be technology which enforces the disclosure
5 of the physical presence of a user agent (resident in a device) to other user agents within the networked environment.
6 Finally, we need to protect users from prejudice, on the grounds of religious beliefs, political affiliations, sexuality, and
7 membership of certain organizations. A user's inability to provide resolutions credentials could result in a
8 significant threat of the fundamental elements of their individual freedom and civil rights. Thus credentials
9 proving the absence of such affiliations/memberships or professions should again be banned from use. Thus, an
10 important condition should be provided as a part of the system's design criteria enforcing the recommendation that
11 users must collectively avoid the use and request of such resolution credentials as part of the user profiles

12 13.17 Additional Applications of Virtual Tags

13 The following is a list of additional types of application specific information which may, for general or specific
14 applications, be usefully deployed through the use of virtual tags.

15 (a) Buyer and seller information - as detailed in the parent issued patent, specific details of what buyers and sellers may
16 be looking to buy or sell respectively may be used to suggest the basis for a potential commercial transaction. The
17 transaction may be large (but not necessarily so, e. g., real-estate, private investment in a small business or public
18 stock). If a physical or on-line interaction with the other party is warranted (e. g., for larger commercial transactions),
19 as is suggested later in the present description, users may identify other users which form the most relevant "match"
20 with their interest. At this point the agents can check for credentials, and then either communicate or calendar a
21 meeting. Similarly, the agents may find the
22 "best" match of users who happened to be physically proximal to the user at that particular time, or at some future
23 time(s)/location(s) which is mutually compatible (similar applications are suggested for matching sales persons with
24 prospective clients, identifying experts to work (individually or collaboratively) on a particular project or problem, to
25 answer a question of an appropriate specialized nature to their area of expert knowledge.) The parent issued patent
26 suggests at a general level these commercial applications. An additional feature described therein involves the use of a
27 decision tree called "Rapid profiling" which can be used in the present context to identify from the most common
28 needs of buyers and "goods" of sellers in general and the known profile data about each buyer and seller individually, a
29 list of questions for each party which most briefly and efficiently determines the complete buyer/seller profile of each
30 party individually.

31 (b) Medical information, such as medical conditions, medical history, active prescriptions, drug reactions, family
32 history, possibly even genetic pre-dispositions (from a genetic profile). Medical insurance information may also be
33 potentially useful for a prospective qualified accessor to be able to readily access in case of an emergency.

34 (c) Social Interests Profile Information -- The parent issued patent also suggests the present application at a general
35 level. For a dating application, users may be matched on the basis of their common interests/preferences and perhaps
36 on the basis of certain information reflecting personality, social or cultural behavior/affinities or psychological
37 attributes. On the other hand, for purposes of meeting casual acquaintances, users may be interested in another user
38 who shares the above characteristics as well as someone who has recently shared similar experiences and/or personal
39 challenges.

40 (d) Physical location information -- Users or advertisers could, for example: a) Query a pseudonymous user database to
41 access profiles that are in close physical proximity and match certain criteria, e.g. live in a certain geographical region,
42 had recently attended a meeting or event (or is planning to attend a particular event) had recently communicated with a
43 friend or associate. In another variation, a user could for example, submit a query pertaining to every user in a
44 particular physical space, e. g., a room, hotel or convention center, e. g., identify all users present here who attended
45 Internet World, 1995.

46 (f) Professional Information/Qualifications - As in the application of matching buyers and sellers, a description of a
47 user's needs or situation with relation to various professional services may be provided as additional data about the

CONFIDENTIAL

1 user. Examples may include: (as above) medical data, professional or business history (as well as legal history) which
2 may be of interest to law firms, accounting firms or various business consultants. Personal, family or emotional
3 difficulties may be of interest to psychologists or family counsellors. Again, users may submit this information as a
4 query for prospective matches, or they may be pseudonymous queries or automatically matched in accordance with
5 criteria specified by the professional. The issued patent application also lists additional applications which could
6 as well be relevant within the usage context of virtual tags.

7 Referral Information - Situations frequently arise in a variety of contexts of human interaction (whether social or
8 professional) in which a user may wish to refer the user they are in contact with to another individual. Often this
9 occurs in a professional services context which a user has a particular need/or other characteristics which make him/her
10 an appropriate match for the services provided by the other party. Or in a business context, often a user will forward a
11 business contact or associate to another colleague who is deemed more appropriate for the particular context and/or
12 scope of business. Likewise, in a personal or social context users may sometimes meet two or more individuals which
13 they observe or perceive share common interests, goals or beliefs or perhaps possess complementary capabilities,
14 knowledge, or characteristics. In each of the above scenarios, virtual tags may provide substantial benefits.

15 For example, the referring user could forward the relevant portion of the profile and identified need of the user to the
16 referring party whose user agent may determine the acceptability of the request and/or the priority with which a
17 communication or meeting could be scheduled (e.g., as could be automatically arranged by/between the two party/
18 calendaring agents). If the referring party's agent is unable to make a decision or priority assessment for scheduling
19 purposes on behalf of the user, the agent could instead try to contact the individual him/herself for assistance (and
20 statistical feedback to the system's data model). In order for these types of referrals to be performed efficiently, the area
21 of expertise required can be specified, and provisions can be made about the type of referrals that a professional will
22 accept.

23 (h) Employer/Employee Information - Users who are seeking employment (actively or unofficially/passively and
24 employers who are seeking employees for specific responsibilities may benefit under the current scheme. An employer
25 may post a description as part of his/her virtual tag (and that associated with his/her company). His/her employees may
26 also have provided ratings and/or annotations which are further descriptive of his/her personality,
27 leadership/management style and skills, work environment which s/he promotes and overall quality. Previous
28 employees for that position (who may also have either provided information about themselves as deemed appropriate)
29 or may also have provided such information as well as pertaining to the position. If not s/he may allow him/herself to
30 be contacted by the prospective candidate (e.g., in exchange for a fee).

31 Access Privileges Information - Users in an organization are frequently given privileged access to certain files within a
32 corporate intranet but not others. Though there are many ways of profiling users according to their level of access
33 privileges to information, the following example is considered: Based upon the position (e.g., responsibilities and
34 tenure with the organization), users may be "classified" into groups according to different levels of access to
35 confidential information. Virtual tags may be used to extend the capability by providing for immediate disclosure of a
36 user's information access privileges to another employee in real-time and in a physical context. Similarly, it is
37 conceivable in either a professional or non-professional environment that users may wish to maintain "secrets" between
38 each other. Access or restriction to/from a secret may be provided in accordance with a particular attribute of a user
39 (e.g., departmental secret, a family secret, personal secret) or by the individual identity of the user, e.g., upon
40 contact/communication is privileged or not to a given secret per the restrictions by its originator).

41 Additional rules could be provided, so that if certain pre-defined events occur, the criteria for release of
42 information associated with the secret are changed. Alternatively, this "information" triggered by an "event" could
43 instead be a message. For example, if a user reads or accesses certain information, meets with a certain colleague or
44 friend send message X. This message could be (for e.g.) a request to perform some task relating to part of that
45 information, a reminder to address certain issue(s) while chatting with the colleague etc. or, per the request of an
46 individual's employer or colleague if a given individual (a sales person) meets with user X send him/her message Y
47 (which may refer to a previous encounter, experience or fact s/he should know pertaining to user X and which may
48 have bearing upon their conversation or professional interaction. (This message could as well be sent by user X as a
49 reminder notification to him/herself at the time of the encounter).

CONFIDENTIAL

1 In another application, the user's access privileges may be used for granting him/her access to restricted physical areas,
2 (thus, the virtual tag effectively may behave like an "electronic door key"). A variation of the technique may be used
3 for granting access to professional meetings, where information access privileges of users must match the anticipated
4 confidentiality parameters for the scheduled meeting. Another application may include the ability to automatically
5 enable access or restrict access, based on payment of fees, and whether or not an individual is a representative or
6 partner of a competing company.

7 (i) Contextual Information -- If user X is a prospective customer, and performed certain on-line or physical behavior
8 that suggests an interest in a product/service, then the advertiser's agent can automatically send an ad (subject to the
9 user's privacy policy). Another condition which could trigger either a notification to the advertiser or the message
10 (automatically) could be temperature or weather conditions (which may affect people's purchasing behavior as well as
11 often driving/traveling activities which affects the agent's scheduling and reminder's follow-up with participants on
12 existing scheduled meetings.

13 13.2 Example: Match buyers to sellers

14 For example, in a mobile sales-force, the system of SDI can first generate an ideal list of prospects, and then help
15 salespeople target products and offers to individuals. We can provide information to salespeople about users, according
16 to the profile of a salesperson (and reputation), and a user's personal terms for data-disclosure. Similarly, a system of
17 SDI in conjunction with the methods taught in co-pending patent << LEIA >> allow automatic detection of salespeople
18 close to users (via an anonymous location market). The market allows matches to be made, but does not reveal
19 anything about a user that the user does not authorize.

20 For example, in a marketing network, with a commission-based sales force. User profiles can also be used to determine
21 user-responses to offers and products (see the methods in patent application << system for personalized ... >>) User
22 profiles are generated from interactions with other vendors, salespeople, and current activities/behaviors. SDI allows
23 profiles to be built from extended interactions across multiple vendors, so long as the user authorizes the same
24 pseudonym for each vendor.

25 Similarly, users can themselves use seller profiles, to decide whether or not to interact with a seller, and vendors can
26 use seller profiles to decide how to allocate new customer prospects to sellers. The profile of a sales-person may show
27 correlations between product sell-rate and the type of product, type of user, that the sales-person interacts with. Initially
28 seller profiles may not be very well related to sales-performance, but instead based on general SDI-style profiling, and
29 wider (eg professional) credentials. Later, as a seller gains experience, profiling can be based on a sales-person's track
30 record (and this will subsume other information).

31 As an extended example, we can also consider a system for handling tasks—e.g. a call center, or a dynamic work-flow
32 system. As new jobs (or calls) arrive into a system, the jobs are automatically routed to the appropriate
33 expertsmachines/sales-people according to the ability of the person to perform the task. The ability is measured by a
34 match between the profile of the person and the type of task, for example on historical performance/feedback, and on
35 other relevant attributes, and also through collaborative filtering techniques. The goal of the system is to allocate tasks
36 in the most efficient manner, across a system of experts (or machines).

37 The problem might also be informational: e.g. find an expert on ancient American civilization for purposes of writing
38 an article, or answering a specific question. Relevant information might include the expert's resume, and the expert's
39 knowledge expertise profile developed from his/her activities in responding to previous queries.

40 Level of expertise might also include the size of projects performed within a particular specialized area, and relevant
41 education qualifications.

42 Example: business-business introductions. Another application domain for privacy-protected match-making, where
43 users are anonymous until an agreement is struck is business to business introductions. For example, it might be useful
44 to automatically identify synergies between businesses (e.g. in infrastructure, technology, or product) -- for the
45 purposes of pursuing an advantageous strategic relationship. If the meeting is between two employees of competing
46 companies, then the system of match-making could also ensure that a meeting is predicated on a particular task that
47 does not cause conflicts with their respective companies.

CONFIDENTIAL

1 We can also use a query-based system to establish a user's relevance to a particular task, or another user—along the
2 lines of the method in patent << >> with the efficient method to profile users.

3 The system of SDI can also be used as a confidential database for the purposes of generating statistics from sensitive
4 data. For example, as a trusted system, manufacturers might be willing to provide information about their productivity,
5 margins, retainment rates, production efficiencies, yields etc. The central SDI server could generate statistics, globally
6 for the manufacturing sector, and then individually for each manufacturer—as it relates to the information provided by
7 other companies. Similarly, it would be possible to use such a system to compare salaries across different universities.
8 While an individual university might be reluctant to reveal information about its pay-scales to other universities, in the
9 aggregate this information is not sensitive—and a survey on salary can be useful to both employers and job candidates.
10 SDI is used to securely calculate statistics, without revealing any information that might compromise the privacy of a
11 single employer.

12 Note: Describe secure virtual database involving matching by expertise and sharing of human resources (e.g. virtual
13 work groups). Each member (a commercial interest) has a certain "resource sharing policy" which defines 1) what
14 entities or types of entities s/he would share resources with. 2) If so, on a per-entity or per-entity type basis, what types
15 of their resources (e.g. type of skilled employee and for what TYPE of outsourced task) would the entity share. It is an
16 obvious extension to look at sharing of code, technology, intellectual property. A major challenge and limiting factor
17 being how well informed SDI, the neutral intermediary can be made aware of the needs/requirements of a company
18 such that it can make evaluations entirely on its own regarding highly confidential materials with which it can
19 accurately predict the basis for a deal WITHOUT disclosing to the prospective recipient what the technology or know-
20 how entails (which could compromise the value of that asset should a deal not eventuate).

21 Further describe an extension/adjunct of the present human/technical resource sharing method whereby corporations
22 (in particular high tech and Internet related) may employ SDI to utilize the above information regarding their human
23 and technology sharing synergies in order to detect and recommend strategic (e.g. equity sharing, merger, acquisition
24 etc.) relationship opportunities between the entities. B to b and even b to c is worth it user behavior combined with
25 text analysis should also provide revealing clues about what types of companies tend to share similar customers and
26 provide similar (complementary or competitive) products and services which may suggest that such synergies are
27 potentially available (Geographic region may also be detected or inferred). Again the disclosure of detailed business
28 information is very helpful and a data release policy defining the parameters for such strategic initiatives may be
29 critical in order to determine what companies may be potential candidates for which initial feelers (of high level
30 information disclosure) would be appropriate to put out to a prospective company to determine mutual interest and/or
31 further basis for expected synergies.

32 Note: Reference FAQ routing scheme in co-pending patent application, also add a section on using customized prices
33 and promotions scheme to assign a price to a task, query response or virtual work group participation (where potential
34 synergies are identified between the parties).

35 13.3 Virtual Tags for Annotation/Information Filtering

36 In this extended application of SDI, we allow users and other third parties to annotate objects (physical and virtual)
37 with meta-information, either to remind themselves about a previous interaction in the future—or as a system of
38 "knowledge learning", where systems of users leave useful information for other users. Information is left in the
39 environment, leaving a trail for other users.

40 For example, the information that is tagged to an object, referred to as a "virtual tag", can contain a pointer to other
41 relevant information, such as a survey of a film by a third party, or the user's own comments/feedback. For example, a
42 restaurant listing could be annotated with meta-information about the quality of the food and service. Such information,
43 when provided by a wide sample of users, can provide robust information about objects. The information that is used
44 by a particular user can be filtered—for example, weighting the opinion of a respected restaurant critic, or weighting
45 the opinion of users with common profiles (when that information is available).

46 Virtual tags can be assigned to objects with physical locations, and the information triggered based on the physical
47 location of a user (using LEIA technology). Virtual tags can be assigned with expiry dates or other time-sensitive
48 information. An individual user might leave an "action item", for example—next time I return to this object (eg. web
49 page/ vendor) be sure to perform this task, enter this query, check this link for new information. As another example,

CONFIDENTIAL

1 after a conversation with an SDI-enabled user it is possible to tag that user with some notes, to remember the
2 conversation the next time the two users meet.

3 The technical innovation that allows this use of virtual tags, in addition to the protection of privacy, is that we allow
4 users to annotate information to objects that they do not directly own through a system that separates virtual tags from
5 the content that is tagged. In particular, tags can be stored (either at the ISP-level proxy, or main SDI server) for
6 associated web pages, and exchanged/retrieved automatically when the object is accessed. The virtual tags can be used
7 in conjunction with target-object profiles that are generated through SDI for web pages (and approved by vendors).

8 Virtual tags can be searched, using relevant terms, locations, or times, and can also contain links to authoritative
9 information, such as audio and/or video.

10 Tags are encrypted, so that only SDI-enabled users can access them. Tags are also associated with the pseudonymous
11 ID of the user that left the information (although they can be anonymous, an associated profile allows more accurate
12 collaborative filtering techniques). Finally, users can leave data-disclosure policies, embedded into tags—to certify the
13 properties of other users necessary to release the information. When tags automatically are time-stamped with location,
14 and time, and other information we allow for this information to be "fuzzed", as disclosed in the section on
15 Randomized Aggregates, to protect a user's identity.

16 In the physical world, implementation of meta-information in a user's physical information, can be viewed via head-up
17 displays, video cam monitors, wearable computing devices, or audio pieces. The information itself can be embedded
18 directly on physical objects, for example on magnetic strips or via visual encoding techniques—or the appropriate
19 information can be accessed from a secure remote database based on the user's physical location (using LEIA location
20 technology); or bar-codes that provide a universal identifier for an object.

21 13.4 Information Personalization Without Vendor cooperation

22 We can also rely purely on information provided (as virtual tags) by users about a web page, or other physical or virtual
23 object, to filter and personalize information that is displayed and recommended to a new user. For example, because
24 virtual tags are stored and accessed independently of the source of the object (e.g. not on the vendors own web server),
25 they can be automatically picked up and associated with a page, even without a vendor that is subscribed to SDI.

26 The user's client-level proxy server can annotate web pages with relevant information as they are generated, for
27 example adding mark-up to relevant information, or suggesting a link that the user should read, with notes that may be
28 useful. Similarly, the database of tags can be used for an advanced web-search.

29 The ability to personalized provide recommendations to a user, can be used in conjunction with a cache engine, that
30 might be located on an ISP server. For users that follow advice, the ranking of recommendations is correlated to the
31 information that the user reads, and therefore pre-caching can be more accurate, because the information about user-
32 recommendations can be made available to the cache engine. Similarly, this is useful for advanced fetching to a the
33 client-level SDI proxy, for example during idle time. Virtual tags associated with web pages are encoded to: (1) prevent
34 a non-affiliated cache engine from using the tags; (2) prevent a competitive infomediary service from taking advantage
35 of the "community information".

36 13.5 Demonstrating Value to Vendors

37 Vendors that subscribe to SDI can provide more attractive offers/products to users, based on information about the
38 wider activities/interests of a user, on other vendor pages, and in the physical world (of course, only to the extent that
39 this information is authorized by the user). Vendors can concatenate information from the client-level proxy, the
40 vendor-level proxy, and the ISP-level proxy. In particular, wider information is vital for a good first-introduction to a
41 user, so that information and products can be made relevant from the start. Similarly, such information (e.g. the user
42 likes this type of music, but has purchased this CD) is very valuable for vendors that sell "content-based" products, for
43 example books, and CDs. Vendors can personalize their service with, for example, a recommender system that
44 interoperates with profile information furnished by SDI.

CONFIDENTIAL

1 We can demonstrate this value experimentally, for example we can offer a vendor a free-trial and present personalized
2 information/advertisements to one group of SDI users, and regular advertisements etc. to another group. The increase in
3 vendor revenue can be estimated from client-level monitoring of the change in purchase volume achieved with well-
4 focused solicitations *on the vendor's own business*.

5 Also, we can monitor the performance of vendors that use SDI technology, and provide measurements/metrics for new
6 and prospective vendors (for example on click-through rates and transaction rates) etc. All of this is possible because
7 the system of SDI has access to a user's client-side proxy machine, and the vendors themselves cannot block this
8 collection of data.

9 NOTE to David -If you agree, it may be clearer to define "infomediary" in conjunction with the iamworthit section as
10 a user-centric SDI which typically incorporates the iamworthit and/or community dollars models and go through
11 and replace iamworthit elsewhere in the spec of "infomediary" use in the title "User "infomediary

12 14. User Infomediary: The iamWorthit Model

13 In this section we describe a key extension of the user-centric SDI model, that provides users with an additional
14 incentive to provide information to vendors/advertising networks, in addition to the benefits from receiving well-
15 targeted information and products. We introduce a new currency, termed "community dollars", that allows
16 vendors/advertisers to compensate users for providing information—but tie the compensation to users making a
17 purchase, so that: (1) users are incentivized to provide information that allows vendors to push relevant
18 advertisements/products; (2) users will also be more likely to make purchases at a site for which they can receive
19 discounts via community dollars; (3) providing users with community dollars will increase the number of hits to a site.

20 14.1 Time-of-purchase Competition

21 When a user specifies this "time-of-purchase" competition option to her SDI client proxy, SDI can automatically
22 provide competitors with information about a user's product or service requirements, and a user's profile, before a user
23 makes a purchase. This will facilitate competition between vendors, and can lead to better prices and offers for users.
24 We allow vendors to opt-out of this scheme, and prevent the system of SDI from informing competitors about
25 offers/purchase-requests. Of course, in this case their user can decide not to visit such sites. The system of time-of-
26 purchase competition enhances vendor competition, and can also help to reduce the costs of entry into a market,
27 because name-recognition becomes less important. New vendors can simply register with the "purchase referral"
28 service, and cherry pick the products that they specialize in.

29 This is a "next-generation" e-commerce service. Current shop-bots, for example "Jungle" at Amazon.com,
30 www.shoptheweb.amazon.com, provides a static comparison shopping service, using static price information. A user
31 can specify a product, and receive price information about the product from different suppliers. However, there is no
32 dynamic competition on price or features. The buyer driven service for flights offered by www.priceline.com is more
33 dynamic, in that a seller is found to match the price that a buyer bids, but does not promote competition between
34 sellers. In fact the sellers can make excess profits from the pricing errors made by users.

35 We can use profile information, and historical transaction information for similar transactions, together with the
36 customer price/promotion algorithm disclosed in co-pending patent << >> to negotiate on a deal with a vendor that will
37 optimize the value to the user. Profiling of vendors, and user transactions, can allow users to avoid making bids that are
38 too high and losing value (airlines in priceline.com can profit from inaccurate user bids).

39 We enable vendors with competing products or services to receive automatic notification when a user is interested in
40 the purchase of a particular product. A vendors can also receive information on the profile of a user, and the offers
41 made by other vendors; and submit counter-offers to a user via the user's SDI-enabled client. The user can then be
42 presented with a final set of offers, before making a purchase decision. In a variation, the user may actually initiate the
43 processing submitting an initial offer to a vendor or a collection of vendors for an item, she/he is interested in (whether
44 in a traditional store-front or even an on-line auction house). IamWorthIt protects the value of a user's profile through

CONFIDENTIAL

1 competition, because instead of using a profile to practice price discrimination and extract excess profits from a user, a
2 vendor must use a profile to offer the most appropriate product to the user at a competitive price.

3 The IamWorthIt model provides an environment in which a market is established that is dictated by users by facilitating
4 competition by/between vendors vis-a-vie the release of information pertaining to offers presented to users of a
5 disclosed profile (or more typically pseudonymous user profiles) by vendors to competitive vendors selling the same
6 products or services. In addition to enabling a more competitive market for electronic commerce, the IamWorthIt
7 model also increases the value of user information, that the user can choose to disclose to vendors to enhance customer
8 targeted promotions and targeted pricing and other on-line marketing strategies, and also to withhold from vendors
9 information in order to elicit optimal offers.

10 For example, the information that a user discloses to a vendor could include its sensitivity to discount offers, customer
11 loyalty with other vendors, value responsiveness, (bargain driven), high quantity purchases (for only those categories
12 which the user makes frequent or large purchases).

13 When a user selects the SDI time-of-purchase competition option the client software monitors the user's web actions
14 for queries and browsing related to purchases. The client identifies other vendors with similar products or services,
15 either using a static "web index" that maintains vendors in particular product domains, or through dynamic profile
16 matches between the target object profile of the web site that the user is currently browsing and target object profiles of
17 the web sites of other SDI-enabled vendors.

18 Purchasables which are similar or the same to another purchasable on other vendor's web sites must be identified in
19 order to determine which notification of a browsing action or purchasing request by the user is of relevant interest to
20 those vendors. This may be achieved by solely relying upon meta data embedded in the pages, or with static
21 information that classifies vendors into product-domain categories, such as a Web directory. Key terms and other
22 attributes associated with the items, automatic classification and clustering techniques applying usage statistics and
23 content (associative, numeric and textual attributes as described on the parent issued patent) may be further deployed as
24 additional techniques for purposes of identifying similarity at the level of the target objects. Classification and
25 clustering techniques can be deployed to identify similarity between vendors at the level of target objects.

26 Vendors are notified, and provided with the ability to access the profile of the user, either with client-level processing
27 or through the release of an anonymous profile to the vendor. Vendors typically will wish to construct offers through a
28 rule-based engine, data-mining techniques, or automatic collaborative filtering techniques, as disclosed in co-pending
29 patent application "System for Automatic Determination of Customized Prices and Promotions" and U.S. Patent
30 #5,754,939, "System for Generation of User Profiles for a System for Customized Electronic Identification of Desirable
31 Objects" as such techniques may be deployed by the vendor directly or via the Secure Data Interchange representing
32 the interests of the Vendors.

33 User profile information may include a temporal profile of the user's present activities, including search terms, recent
34 page navigations, what pages is the user observing presently (and the profile of this page) or even his/her present
35 physical location as well as the general user profile (any portion of the above particularly the latter may of course be
36 withheld from the vendor). In the preferred implementation, vendors are also provided with a (client or web-based)
37 rules interface which enables the vendors to input pre-stated rules with which the system may solicit and respond to
38 competitive offers automatically. If pre-stated rules are used to automatically respond to a notification with a
39 competitive offer, the nature and degree of discount is typically determined in accordance with the nature and degree of
40 the original or previous offer and/or the user profile as disclosed by the client-level proxy/server to that vendor. In lieu
41 of manually entered rules, co-pending patent application entitled "System for Customized Prices and Promotions" or
42 another similar algorithmic methodology may be used as an aid by the vendor in order to automatically determine a
43 competitive offer (or subsequent responses thereto). These techniques can also be used in conjunction with a data-
44 mining interface, in which predictive metrics as to selection, price and promotional type, may be determined in relation
45 to the individual user or specific relevant user profile attributes for example, in accordance with a data analysis expert
46 of the vendor (or representing the vendor via SDI) analyzing randomized versions of user profiles and randomized
47 aggregate statistics.

48 A variation of the system is further disclosed which was intended (according to the co-pending specification) as an
49 electronic assistant to telemarketers and other sales persons to determine offers and counter-offers which are
50 automatically generated in response to (for example) rejections of the previous offer as well as counter offers
51 by the user. This dynamic system was originally designed for salespersons to optimize the expected profit from each

CONFIDENTIAL

customer (in view of the general user profile and the offer user responses up to that point in the negotiation). As such, this technique could be readily extended to the current application in which the previous offers up to that point may instead originate from other vendors (instead of a single one), thus the system responses may be affected by the user profile as well as the offer response pairs up to that point in the negotiation process.

It is likely that vendors will not compete on price alone, but rather through added-value services such as offering loyalty bonuses, cross-sells, and two-for-one offer and added features. Vendors will choose this mode of selling to prevent simple price-comparison at the client conversely vendors may attempt to eliminate the features in order to create the perception of a better deal through marginal price reductions. Therefore the client will receive offers from multiple vendors, and after initial filtering of dominated offers, present a choice set to the user. Furthermore, we can allow vendors to offer payment to a client in return for displaying an offer to the user, and vendors can also bid for space on the user's web portal which is often represented as a profile associated with a pseudonym conjunction with a description of the ad space. The purchasing decisions of the user may be performed by an electronic representative of the user's wishes (as "user agent") implementing the techniques of pricing/promotion selection algorithms completely autonomously on behalf of the user. However, the best offer can only be presented to a user to the extent that the SDI client-level software understands a user's model of "value", and can make appropriate tradeoffs between product features and price (as implicitly inferred by the system through the above suggested techniques or explicitly stated by the user in advance). Nonetheless, this is a hard problem, and we expect that the user will often need to make a final product choice decision.

The collaborative filtering techniques described in patent "System for Automatic Determination of Customized Objects and Promotions", can allow a user's client-level proxy server, termed the user agent in this section, to automatically analyze offers. The system can also be used to send initial offers to vendors, on the basis of historical information about the transactions that have been performed between other users and the vendor. Offers can (of course) be sent to a vendor and its competitors. Finally, after offers that are received from vendors are pre-screened, they can be automatically ranked for value—using a combined quality and price metric (again judged within a collaborative filtering framework). The goal is to leverage the database of other offers that have been accepted by users in the past, and form a model of vendors, to determine whether or not a user has received good offers. (i.e. we can exchange information within the system of Secure Data Interchange, and making more information available increases the efficiency of the market). Offers can be filtered and presented to a user in rank order.

Buyers might also form "buyer coalitions", on the basis of automatically detected synergies between their requests. This can give buyers more leverage in negotiation with a vendor.

If the user so desires, the client-level proxy can also automatically notify these vendors if/when a particular offer is about to be accepted by the user. For example, a time delay response in the client-level proxy actually processing the order requests could allow vendors a final opportunity to present another competitive offer to the user. In another less optimal variation, vendors are notified only upon the user agreeing to accept an initial offer received.

As an additional service to users the SDI-level proxy server can perform analysis on the offers that a user receives, through comparison with offers that have been received by other user with the best offer that has been received by any user for the same product, and with the typical offer received by a user with a similar profile to the user. This can be useful to a user because it will allow the user to reject all offers if they are non-competitive. The SDI-level proxy could also automatically identify for users the profile attributes that promote good offers, and the profile attributes that promote bad offers, as an informational service to enable users to gain better offers in the future, either through only revealing certain information or changing behavior to attain favorable profiles.

As an additional service to vendors SDI can provide enhanced profile information, aggregated from other vendors, to enable vendors to provide better focused offers than can be provided on the basis of the profile information directly associated with the pseudonym of a user. Certain portions of the user profile data that is unavailable for direct collection by the vendor (such as information that is collected on other sites including, in particular, competitive vendor sites) may reveal important information which enables the vendor to better target that user. As such the secure data interchange representing the collective users may aggregate, analyze and sell this data to the vendor.

14.2 Community Dollars

CONFIDENTIAL

- 1 We allow users to receive compensation for providing personal data to vendors, information that has value to vendors
2 because it allows information to be focused (for example relevant ads can be displayed to a user, based on his/her
3 profile). The system of iamworthit credits users for information, and provides users with direct incentives to reveal
4 profile information to vendors.
- 5 A vendor can sign up with iamworthit.com and agree to provide only the most restrictive type of community dollars,
6 that can be spent at that vendors site.
- 7 Community dollars are the currency that vendors provide in return for the right to provide focused information to users.
8 Dollars can be general (e.g. for a network of vendors), or very tightly focused (e.g. for a particular product, at a
9 particular time). The user-centric infomediary acts as a broker, matching users and vendors. Another key role of the
10 infomediary (e.g. the portal) is to protect the user from information saturation by controlling the flow of solicitations.
11 (i.e. restrict the number of ads. that a user sees)
- 12 We can use meta-tags to restrict the way that community dollars can be spent. The tag associates the dollar, but the
13 dollar is released within the system of blinded signatures (Chaum) so that a user that collects dollars over many
14 transactions with different vendors can spend the dollars without compromising his/her private information about
15 pseudonyms.
- 16 ** DP. Need to disclose a system for aggregating dollars on the client-level proxy, for dollars collected under different
17 pseudonyms.
- 18 Dollars can be restricted to a number of vendors, and also restricted in additional ways—i.e. they can only be sent if the
19 user visits the site through a particular portal, cannot be redeemed at a competitor, are worth a bonus if redeemed with
20 certain vendors, etc.

21 14.3 The iamworthit community advertising dollars model

- 22 The primary objective of the iamWorthIt model is to create a market for information about users. We allow vendors to
23 pay in "community dollars" for adverts, dollars that can only be spent at that vendor (with the host site of the advert
24 receiving a share of the profits). This provides vendors with the ability to gain long-term customers. Furthermore, so
25 long as the user agrees to receive advertising from his/her iamworthit subscription offer, community dollars can be
26 replenished at the rate at which advertisers are willing to pay for impressions. This provides users with an incentive to
27 spend at the vendor's site, because the vendor can monitor (pseudonymously) the user's that are sensitive to discounts
28 and other special offers (that are delivered as community dollars).
- 29 ** DP. Be sure to disclose the method that vendors can track the user that receives community dollars without
30 compromising a user's pseudonymity.
- 31 The main mode of the community dollars advertising model allows vendors to advertise for free, but provide
32 community dollars to users, that can be spent at some later time. The cost of advertising can be linked to the success of
33 advertising. Moreover, the vendor can direct offers and adverts to particular user profiles. The hosting web page
34 receives a share of the vendor's revenue that comes from transactions involving community dollars. The dollars can
35 represent "stored value", such as bonus points, that can be applied to special discounts for offers which are delivered
36 via digital coupons and/or as "straight value" which could be converted directly to purchases thus are equivalent to real
37 dollars at the point of transaction.
- 38 The community dollars can be "credits" that can be redeemed as real cash, credits towards discounts, and can be spent
39 across a suite of sites, or limited to one site. The co-pending patent application entitled "System for the Automatic
40 Determination of Customized Prices and Promotions" describes a comprehensive scheme which may be implemented
41 in either on-line or off-line commerce environments. The system enables vendors to deliver a digital message in the
42 form of a promise to a user (typically on encrypted form for purposes of targeting a user specifically). This promise is
43 typically a discount for a product, set of products (or all products in stock) or may even include entitlement to special
44 privileges for that user, thus it is termed a "digital coupon". The community dollars can represent special discounts for
45 a user.
- 46 The user receives a financial incentive for receiving well-targeted solicitations, while preserving user privacy within the
47 SDI system. The vendors support the community dollars through advertising revenues and increased sales volume. We

CONFIDENTIAL

1 can also provide the vendor through which the user first subscribes a special "first screen" right, that allows the vendor
2 to provide a user with his/her first impression as soon as s/he logs on.

3 In one variation all community dollars collected by a user must be spent back at the vendor site at which they originally
4 subscribed (and also the site that hosts the adverts of other vendors). A user can spend the dollars with any vendors that
5 are site partners of the original site. This provides the vendor an incentive to accept and promote the community dollars
6 concept.

7 NOTE: DAVID Since the community dollars concept is unquestionably the most critically important piece of iReactor
8 technology I want to be ABSOLUTELY POSITIVE it is comprehensibly disclosed technically.

9 The value of providing a user with targeted solicitations is estimated at approximately \$300 to \$500 per year (based
10 upon \$120 per 1000 targeted impressions at approximately 25 impressions per day). Given these significant benefits, a
11 vendor can provide a user with a significant discount (in the form of community dollars). Vendors benefit from
12 increased sales volumes.

13 When the price of items is less than the value of dollars, the vendor can limit the amount of discount that is available on
14 any single product, or only allow community dollars to be applied towards customer discounts (which may
15 nevertheless be quite substantial).

16 14.4 Preferred Implementation

17 In the preferred implementation we use an "electronic cash" infrastructure for the community dollar system. A user's
18 SDI-enabled client-level proxy stores dollars that the user receives securely. Dollars are anonymous and non-traceable,
19 so that the user can maintain a single "bank" of dollars, and aggregate dollars collected across pseudonyms for a single
20 purchase, so long as the purchase satisfies the constraints on the dollars. Each dollar is created using Chaum's blinded
21 signature technique, and also signed with the conditions on its use.

22 ** DP. Need to disclose a technique to allow transfer of dollars across a user's pseudonyms.

23 This scheme allows vendors to monitor the offers that users respond to, because when a user presents a community
24 dollar—the dollar can be validated to indicate the type of discount that it is, even if the identity of the dollar (i.e. the
25 serial number) is untraceable. SDI provides vendors with guarantees that users have once-in-a-lifetime pseudonyms, so
26 redeeming a voucher of a particular type that is redeemable only at vendor V and was issued by vendor V allows
27 vendor V to be sure that the voucher was issued under the same pseudonym, and has not been transferred to another of
28 the user's pseudonyms.

29 In an alternative architecture, the ISP-level SDI proxy, or the web-host for the advertising service, can maintain
30 community dollar "debit" accounts for each user. This is more limited, because it does not allow users to transfer
31 dollars between pseudonyms without compromising privacy (revealing a portfolio of pseudonyms). However, in a
32 scheme where advertisers require that agents have once-in-a-lifetime pseudonyms, and only release community dollars
33 to be redeemed at their own site, this is not limiting. Both of these approaches are useful for "community dollar-
34 enabling" numerous or all sites.

35 A vendor that allows community dollars to be spent does not need to implement a special community dollars/discounts
36 program. The user can also be issued a special debit account dedicated to community dollars, that permits
37 pseudonymous transactions without revealing a user's portfolio of pseudonyms.

38 ** DP. disclose how community dollars can be integrated within a general privacy-protected SDI banking system. e.g.
39 the ISP level SDI proxy allows a user to transact via the general SDI credit card account.

40 A portal site that hosts advertisers and users that subscribe to IamWorthIt can mandate that all community dollars are to
41 be spent at sites that advertise on the portal site, and also only when the sites are accessed via the portal site. This
42 technique will increase portal traffic. Portals can be expected to compete in terms of: (a) the fraction of advertising
43 revenue that is turned over to users, in return for receiving profile information from users; (b) the level of advertising
44 that users are exposed to; © the nature of the community dollars "package", i.e. what vendors can the dollars be used at
45 etc. This can be useful to attract niche customers, that have common outlooks, interests, and business needs. The
46 primary goal of the portal is to drive traffic through the portal.

CONFIDENTIAL

1 ** DP need a technical solution to ensure that \$\$s can only be redeemed if transactions are done via the portal.

2 14.5 Commercial Variations

3 (i) Vendor coalitions

4 Vendors may choose to form coalitions, to allow users to spend community dollars at any "partner" site. Vendors
5 that have similar user bases can be automatically identified using collaborative filtering. (i.e. determining similarity
6 with the present vendor, from the aggregate vendor preferences of a given vendor's subscribers). Also, these
7 resulting metrics could incorporate predicted online spending by each user at each site. This could help to narrow
8 the selection of sites the vendor wishes to partner with and/or the selection of these partner sites could be
9 determined and presented to the user to even further narrow the selection for each user. All vendors in a coalition
10 advertise, and provide cross-links and up-links to other vendors.

11 The coalition model is good for users, that are more likely to find products that they want. Vendors can share the
12 risk of advertising, since dollars provided to one user by a particular vendor can be redeemed at another vendor.
13 Advertising and community dollars increases sales volume at all vendors in the coalition.

14 We can even allow vendors to form dynamic and virtual coalitions within SDI, with a potentially unique coalition
15 of vendors for each user. The coalition may consist of an optimal pool of vendors, as determined by SDI
16 collaborative filtering techniques. The goal in this model is to provide users with a particular "brand" of
17 community dollars.

18 We can allow each vendor to retain an exclusive right to advertise to each user; and also develop a portal for the
19 coalition—that gives advertising prominence to coalition members. Portals will be expected to aggressively
20 promote community dollars. Users that collect community dollars become loyal return visitors to the portal and its
21 associated vendors. In the case the vendors do not generate the same value we can provide community dollars in
22 proportion to the value that a vendor contributes to a coalition.

23 We can also provide targeted advertisements for the vendors at the portal, using the user profile to focus ads. The
24 categories and links at a portal (that might include a search engine) can be re-prioritized (highlighted and/or re-
25 ranked) in accordance with the user's preferences (as described above), and to favor subscribing vendors.

26 Vendors pay the portal site to advertise, and the portal provides community dollars to users in return for privacy-
27 protected profile information. This model does not provide incentives for the portal to provide well-targeted
28 adverts, because there is no direct link between a portal's revenue stream and the vendors' sales volumes.

29 A portal with community dollars that can only be spent under a single pseudonym at its partner sites also provides
30 an incentive to users to interact under a single pseudonym—which in turn allows a portal to profile users across its
31 complete vendor partner network. Users will access many sites with the same pseudonym. The system of SDI
32 allows vendors to leverage the shared profile information as users browse web pages and products.

33 We can also lock users into a single portal—and a single coalition of vendors—with community dollars that
34 "decay" over time, and must be continually replenished. In this way a user cannot pick and choose different
35 portals, and different community dollars, but can benefit mainly from high web-browsing volume through a single
36 portal. The value to vendors in terms of consumer lock-in can be considerable.

37 For example, a coalition of vendors can join to allow a user unlimited access over all affiliate vendors. The
38 program can be sold through the existing marketing channels of each vendor, as well as through a portal directory
39 of sites for those vendors. Vendors that join can be required to promote the program through their own marketing
40 channels. Additionally, perhaps vendors are selected to cover exclusive physical regions (e.g. in the case of a set of
41 ski resorts), or exclusive product categories (e.g. in the case of on-line vendors). Vendors can provide a community
42 dollar-for-real dollar exchange, in return for becoming part of a vendor network. Alternatively, perhaps vendors
43 provide an up-front fee, that can be recovered via dollars spent by users at their own site. Each vendor is obligated
44 to sell the partner network community dollars, but is not necessarily required to promote the other community
45 dollar vendors.

46 (ii) Enabling "Transaction-based" Revenue-sharing

CONFIDENTIAL

1 In a second mode, the vendors provide discounts to users, and advertising is primarily "free". The only time that
2 vendors incur a cost for providing a user with an impression is when they achieve a sale — because then the portal
3 site receives a cut of the transaction price. The vendors provide users with community dollars directly. The dollars,
4 which are stored at the portal site, allow user-spending to be tracked. This allows the portal to monitor when a sale
5 occurs, not just a hit on a banner ad. With transaction-based revenue of this kind, personalization is critical. In this
6 model the portal with give prominence to adverts from successful sites.

7 A portal site may forgo payment from a vendor in exchange for the increased click-through from a strong network
8 of community-dollar enabled vendors. Value is credited directly to users for future redemption at that particular
9 vendor's site. The community dollars provided to a user can be restricted, such that a user can only redeem dollars
10 if the s/he maintains enough visits to the portal site.

11 Vendors can offer discounts on their own products directly, instead of providing the portal with money for
12 advertising. The vendor only pays to the extent that its advertisements are well-targeted. The vendor could also
13 request special ad priority. A vendor that presents advertisements to a user offers the user discounted promotional
14 offers for products offered by partner vendors, in exchange for subscribing to iamworthit and receiving targeted
15 impressions. These offers are in lieu of community dollars, and can be provided by partner vendors — maybe in
16 exchange for a right to a number of ad deliveries for the vendor's own advertising purposes.

17 We can also require that users are automatically routed through a portal when accessing any partner vendor
18 directly. The portal (and therefore the coalition of vendors) then receives exposure each time the user clicks on an
19 ad (or link) to that vendor.

20 The portal may also provide benefits (e.g. additional advertising prominence) for sites which are responsible for
21 driving traffic through the portal. Community dollars can be provided whenever the user accesses a site from the
22 portal. Portals can offer free advertising to e-commerce sites (forgoing advertising fees). The portal provides
23 discounts to users that purchase a product following a link provided at the portal.

24 A user receives the discount by validating a purchase with the portal, and the site agrees to provide the portal with
25 a share of revenue whenever the user cashes in community dollars in this way (we do not rely on HTTP refer
26 mechanism because that can be blocked and falsified. Furthermore, we do not rely on URL+extension
27 correspondences, also not secure — instead rely on providing user's with incentives, and monitoring users that have
28 earned community dollars.)

29 (iii) Delivering per-impression dollars

30 We can deliver community dollars on a per-impression basis, with vendors competing to offer users high values
31 for being able to present an advert. The existing collaborative-filtering engine at a user's SDI client-level proxy can
32 filter ads, and select appropriate offers, using community dollars as just another measure of the usefulness of a
33 message. This is an alternative to providing dollars on a one-off (or even yearly) basis, for consumption via the
34 vendor's site that the user subscribes to the service.

35 A hosting site can take a fraction of any dollars provided to a user. Alternatively, a site can convert the value into
36 the community dollars to provide to the user, possibly at a preferable basis. The portal might also wish to convert
37 its commission to credits for the user at any one of its partner vendors, with the stipulation that the user must
38 access those sites via the portal in order to be able to redeem the credits.

39 14.6 Providing Loyalty Bonuses

40 We can use the client-side SDI proxy to provide vendors with "loyalty guarantees", that are credentials to verify that
41 the user has executed no transactions with any competitor, under any of its pseudonyms. The client-side SDI proxy is in
42 a unique position to be able to implement this monitoring, because no other system knows a user's portfolio of
43 pseudonyms. The user can present its digital credential when visiting a vendor's site.

44 A vendor may wish to provide loyalty dollar credit; for example, it would be possible for vendors to offer user's credits
45 if the user is a 100% loyal customer i.e. that she/he did not (over a specified period) do his/her purchases at the site of
46 any competitor. For example, certain types of high value customers could be given considerable value in the form of
47 credits or discounts as a result of demonstrated vendor loyalty.

CONFIDENTIAL

- 1 A competitor could be determined by a simple known vendor list, e.g. through Yahoo or by matching similar products
- 2 or service descriptions on sites across the web (e.g. using web crawler technology).
- 3 The credential can be time-stamped, to prove loyalty. It does not reveal any information about the user's other
- 4 pseudonyms to a vendor, because many pseudonyms will exist that have not made any purchases from a competitor.
- 5 Upon accessing the vendor's site, this credential may be presented to the vendor. One criteria for the above benefits
- 6 could be that the user may visit a competitor site, and engage in interactions; however s/he should not transact with
- 7 that vendor.

8 14.7 Off-line Variations

- 9 The community dollars may be stored on a portable smart card carried by the user, in which the community currency
- 10 was originally loaded by the user's own computing device, or alternatively a kiosk or cash register controlled by the
- 11 vendor or another third party whereby redemption may occur at any location where the subscribed vendor possesses a
- 12 smart card reader. In another variation, the community dollars may be coded into a form which is bar code reader
- 13 enabled and distributed to the user electronically or potentially if used in conjunction with a traditional loyalty points
- 14 program, additionally printed for the user at the vendor's physical location (such as point of sale or kiosk), applied in
- 15 conjunction with purchases at the vendors physical location. At which point typically a new coupon is reprinted
- 16 containing the updated secure information pertaining to the user's community dollar and/or loyalty points account.
- 17 In another variation, a promotion for a yearly allowance of community dollars could be printed as an advertising offer
- 18 on a magazine coupon, newspaper insertion or direct mail piece which could contain a unique URL (typically the actual
- 19 URL for the iamworthit community dollars subscription site with a unique post script as the character string ("code")
- 20 identifying that particular vendor and/or that promotion) from which the user could subscribe to iamworthit, wherein
- 21 the unique URL acts as an identifier for that particular vendor's promotional piece from which the user originally
- 22 received the offer for his/her own community dollars promotion.
- 23 ** DP. Careful, this could be used to breach a user's privacy, i.e. link a pseudonym and profile with a user's true
- 24 identity.
- 25 Within SDI we could send physical solicitations to users, and allow users to access promotions pseudonymously. SDI
- 26 can target a selected audience for each vendor.
- 27 ** i.e. we need to user to be able to subscribe without a link made to real identity.
- 28 Example: An iamworthit card in accordance with the pseudonymous payment methods described above, such a card
- 29 could be a direct extension of SDI into the off line environment. Users could use this card as an identifier such that
- 30 when they travel physically from vendor to vendor, their profile data can be readily identified where data pertaining to
- 31 their own behavior and policy (depending on their data release potentially part of the vendor's user profile data) is
- 32 retrieved.
- 33 If a smart card is used this user profile data may not have to be remotely retrieved but may be stored on local memory
- 34 on the card itself. In one novel variation, a card is done away with completely by virtue of revolutionary technological
- 35 breakthroughs in being able to instantly and positively identify users biometrically using ISP's scanning techniques
- 36 (which may in a variation be further combined with facial recognition techniques). Many vendors will wish to utilize
- 37 user profile data in order to deliver promotions targeted discounts and promotions (see pending patent "System for
- 38 Customized Prices and Promotions").
- 39 The co-pending application entitled "Location Enhanced Information Architecture" (LEIA) describes an integrated
- 40 advertising delivery platform which selectively targets user personalized advertising based upon both the user's
- 41 personal profile and the present location of the user which may suggest appropriate ads from vendors which are local to
- 42 the user, wherein user identifiers (UID's) which could include any of the above identification media provide the
- 43 essential elements for this user targeting platform.
- 44 For example, at a bookstore, we can recommend isles and particular books; at a supermarket, can play music
- 45 preferences; smart-radio, play appropriate channels in a cab based on target object profiles (as meta-data). As
- 46 suggested in issued patent "System for Broadcast of and access to Video and Other Data Using Customer Profiles" the
- 47 appearance of relevant selections can be continuously scanned for, dynamically selected and presented to the user in the
- 48 form of "Virtual radio station". Such a system can also be linked to a service for making an instant purchase, or linked

CONFIDENTIAL

1 to a database (in conjunction with LEIA) to recommend where a user should physically go to make a purchase. For
2 example, music selections that the user is presently listening to may be ordered.

3 Also, it is possible to provide advertising which is targeted to a user by automatically recognizing pre-existing
4 commercials and replacing them with targeted counterparts. This can be done through the identification of previously
5 played commercials for example, commercials that have been manually identified and classified. Upon recognition,
6 targeted commercials (including those which are targeted by user location in accordance with LEIA) may be inserted
7 into these spots, and delivered and/or pre-cached through cellular, satellite or radio communications.

8 At a public phone we can identify a user with his/her calling card, and deliver targeted advertisements, via the public
9 telephone readout or delivering the targeted ads as audio messages in which server software at the phone switch (an ISP
10 level proxy) recognizes if/when the user is put on hold and delivers audio and/or audio/video advertising to the user
11 accordingly.

12 Targeted discounts and advertisements can be delivered at kiosks, for example using a credit card/smart card/other ID
13 method (e.g. biometric...). Similarly, we can use credit cards to deliver targeted print advertisements on the backs of
14 purchase receipts, e.g., supermarkets or fueling pumps or, alternatively, on a sheet dedicated for a advertising purposes
15 conjunction with public copiers or printers or in another variation, on the cover sheet of incoming faxes which are sent
16 to the user's fax machine or in which the user is otherwise identified automatically from the recipient's name field on
17 the cover sheet.

18 Similarly, we can deliver targeted advertising and other information through cable TV systems, as described in the
19 issued parent patent application entitled "Broadcast of and access to Video and other data users customer profiles", and
20 co-pending application entitled "Broadcast & System for reduced memory terminals broadly address the use of cable
21 systems as an interactive medium (in a bi-directional network architecture) for purposes of delivering targeted
22 advertising targeted advertising and other information to the consumer based on user profiles". In this system customer
23 behavioral data is collected at the digital set top and the upstream channel enables these profiles to be processed at the
24 lead end server. These detailed profiles may then be subsequently transmitted down and stored at the level of the
25 individual set top. The cable environment is a two way interactive medium. The bandwidth allocation is inherently
26 asymmetric. Separate channels can push parallel adverts, which are selected at the set-top-box according to a user's
27 profile. Each channel can have associated meta-data to allow matching at the set-top-box. As an alternative variation,
28 full motion advertisements may be down loaded in the form of applets to the digital set top box and displayed to the
29 user in similar fashion as described above.

30 14.8 Example: An on-line Gaming Site

31 Consider an on-line gaming site that has a network of affiliated vendors, that do not pay to advertise, but provide
32 community dollars that can be spent either at the vendors--or at the casino. When users lose money at the casino the
33 casino receives real dollars from vendors. The casino is one possible outlet for spending dollars—and a vendor only
34 pays the casino (the host of its ads) if a user chooses to gamble on the site, and loses its dollars.

35 The gaming site becomes a portal, with links to partner vendors. Each vendor offers the user community dollars, that
36 can only be spent back at that site or at the casino. However, the number of dollars which can be spent at the store is
37 substantially less than the number of community dollars which can be spent at the casino. IF the user accepts the
38 credits, whenever the user accesses the URL to the store he/she is either automatically routed first to the casino portal
39 or to the vendor site whereby a prominent banner is displayed which is displayed to that particular user from which the
40 user can conveniently engage in a casino gaming session.

41 If the user loses a substantial amount of community dollars he/she may regain the lost credits by spending a specified
42 amount (in real dollars) at a partner vendor. This provides a safety-net for users. The cost to the vendor is the cost of
43 the dollars that the user lost at the casino, and the cost of replenishing the user's community dollars (which can be used
44 for further gambling). However, the vendor makes a sale—so the vendor is happy so long as the dollar value is a
45 reasonable discount for the sale.

46 Clearly, the casino gains substantially through the redemption of these community dollars. The vendors can make an
47 agreement with the casino where they only compensate a fraction of community dollars. A percentage of a user's

CONFIDENTIAL

1 profits at a casino are paid in community dollars, another percentage can be paid in real dollars. Perhaps the casino can
2 also provide vendors with a revenue share.

3 14.9 Implementation Details

4 We allow for community dollars that are restricted to particular products, and customized for an individual user. The
5 dollar-object can contain two parts: the first part is readable to the user, and indicates the nature and the amount of the
6 discounts to which the credits can be applied. The second part of the message is encrypted, and accessible only to the
7 vendor, and is signed by the vendor to prevent any form of alteration. The information can contain the dollar credit to
8 the user, the terms and conditions of the community dollars, a dollar amount, the pseudonym ID of the user, an
9 expiration date, the terms and conditions of discounts and special offers to which community dollars may be applied in
10 combination with a partial cash transaction.

11 The vendor must check that it has not previously redeemed any piece of community currency with the same identifier;
12 the identity of the user is correct; the date; the terms and conditions. Some vendors may allow redemption of the
13 community dollars at other vendors' sites.

14 ** DP: See the paper in 3rd USENIX 1998, by Doug Tygar for technical solution.

15 14.10. Privacy Enhanced Marketing Applications

16 There are a number of commercial applications for a system that allows vendors to contact privacy protected lists of
17 users in the pseudonymous user data base, wherein the contact, interaction, and business relationship with the vendor
18 occurs under terms of complete user pseudonymity. Or wherein the user remains pseudonymous in his/her search and
19 access to qualified professional services. In accordance with the parent patent application the pseudonymous
20 communication may be either email, real-time text communications, voice (such as the pseudonymous telephony or
21 Internet telephony).

22 For example:

23 (i) Financial Advice and Financial Planning Services

24 Often users are quite sensitive about the confidentiality of the release of this type of information related to personal
25 financial matters and particularly with certain matters (and perhaps in general) prefer that their financial advisors
26 were unaware of their true identities. Similarly, investment advice or sales communications by stock brokers are
27 another application where similar user information is typically disclosed.

28 (ii) Insurance Agents & Brokers

29 For many types of insurance, (e.g. health, life, casualty) personally sensitive information is disclosed by users to
30 their agents and brokers. Initially, before insurance services are purchased, it is possible that useful detailed quotes
31 and/or insurance advice could be provided to a user pseudonymously.

32 (iii) Legal advisors

33 There are a variety of legal disciplines in which the associated legal services delve into highly sensitive personal
34 information (e. g., bankruptcy law, divorce law, criminal law, etc.) Many lawyers also offer to first-time
35 prospective clients a free consult in which such a privacy-enhanced communications system could be initially
36 beneficial to the parties.

37 (iv) Family Counseling and Psychological Counseling

38 The parent patent application also suggests these applications which often involve the exchange of highly
39 confidential personal information.

CONFIDENTIAL

1 14.11 Business Models Leveraging Community Dollars

- 2 Free or discounted retail products with "niche" partners in each category Free dial-up ISP (as an independent ISP or a
- 3 service to jointly promote free access with ISPs)
- 4 Free Cable and ISP service, Free pay-per-view (note that viewing patterns and the associated content could provide
- 5 additional valuable user profile information)
- 6 Free phone service (e.g. advertise subscription service on screen phones or audio ads from pay phones)
- 7 Free prepaid calling card
- 8 Free print media subscriptions (magazines, newspapers)
- 9 Free book clubs
- 10 Offer any combination of the above with "deep discounts" for each (this can involve \$350 community dollars per user
- 11 or it may simply involve certain purchasing limitations per customer). Each vertical niche partner in exchange gets
- 12 exclusivity within their own respective niches to target advertise to those users (e.g. retailers).
- 13 Free access to sporting events.
- 14 Free credit for casinos
- 15 Free lottery tickets
- 16 Free charity donations
- 17 Discounted hotel lodging
- 18 Monetary credit to a credit or debit card (either an iamworthit branded card or provided as a partnership with the card
- 19 companies.
- 20 Monetary credit to a diner's club
- 21 Free subscriptions plus credit to retail buyer's clubs (on-line or off-line)
- 22 Credit or discounts for book clubs
- 23 Free musical concerts, or theater presentations, movies or
- 24 access to arcade entertainment
- 25 Free access to amusement parks or theme parks
- 26 Free golf season passes
- 27 Free commission fees for stock trading
- 28 Free commission fees for travel booking (if implemented for on-line users would be less compelled to search for travel
- 29 information on-line though go off-line to make their bookings).
- 30 Providing a substantial credit in a user's electronic wallet in exchange for their downloading the electronic wallet to
- 31 their client.
- 32 Free personal home pages (which community dollars could subsidize a high quality site).

33 14.12 Iamworthit Marketing Strategies

34 A number of marketing strategies are worthy to note. These include the following:

- 35 1. Allow the ISP to promote free Internet access through traditional means of Web advertising (impressions on ad
- 36 servers) Web advertising whereby this advertising is effectively traded, directly for (e.g. pop-up) advertising on
- 37 iamworthit. Furthermore, the ad server would be able to recognize through the associated domain names, the users
- 38 which are coming from a competitor ISP. So long as that ISP is not a partner of iamworthit, the associated user
- 39 would be selectively targeted with an offer of this sort "free Internet access" by subscribing to iamworthit".
- 40 Smaller ISP's would be particularly compelled by such offers to their direct target prospects. This is because they
- 41 are operating on a "thin margin". Furthermore, both they and their small regional counterparts would be
- 42 particularly vulnerable to this type of advertising by regional competitors from the same geographical area, during
- 43 specified period of months of initial usage of the service, the share of profit due iamworthit could instead be
- 44 committed to purchase additional advertising for the Internet service provider (or the balance of this profit traded-
- 45 out in the form of additional advertising through the ad server partner).
- 46 2. The ad server partner could further become an exclusive partner of iamworthit on the following commercial
- 47 venture:
- 48 Relationships as established with on-line merchants and other e-commerce sites. The vendor actively promote an
- 49 offer to their customers through both off-line media (using a URL unique to that vendor) and on-line advertising
- 50 through the ad delivery partner. The offer may say (as an example), "receive three hundred dollars in purchasing
- 51 credit at Books a Million in exchange for subscribing to iamworthit (or receive five hundred dollars worth of
- 52 discount credits at Books a Million. In addition, as a further benefit to the vendor iamworthit could trade its own

CONFIDENTIAL

- 1 advertising impressions with impressions on the ad server for the current offer (in order to reach a different base of
2 users which are not currently subscribed). A particularly compelling industry for this application is on line travel
3 inasmuch as a plaguing problem to this industry is the fact that many users use travel sites as an informational
4 resource on available travel deals and packages, but ultimately book their trips directly through a travel agent (thus
5 cutting out the travel site). A three hundred dollar a year travel credit would be a compelling incentive to many
6 users to modify their current travel booking habits. Affiliate networks are also an ideal channel for these types of
7 promotions because affiliate sites agree to participate (typically) purely based upon the degree of the profit sharing
8 opportunity (which would be significantly larger than most types of transaction - based affiliate advertising).
- 9 3. Marketing Network Concept to Sell iamworthit - sites which offer a community dollars promotion could, upon the
10 users subscribing to iamworthit, additionally offer the user with a down-loadable client based software which
11 provides a small promotion in conjunction with a link to iamworthit. Each time a recipient of the email subscribes
12 to iamworthit, a percentage of the value of that customer is credited back to the user in the form of community
13 dollars. Each subscriber resulting from the current subscriber's email (though reduced) provides an additional
14 credit to the original subscriber in accordance with the marketing network business model. If the site originally
15 delivering the promotion is not an e-commerce site, a percentage of advertising revenues resulting from the
16 subscriber (and potentially all resulting subscribers) could be used. It could be applied in the form of iamworthit
17 advertising (or exchanged) for advertising in an ad server.
- 18 4. Free Community-based Content (e.g. broadband over the Web) - As an alternative to the community dollars
19 scheme, particularly if a site does not have a principle focus on e-commerce, an attractive proposition to Web-sites
20 could be the creation of premium content which is free to iamworthit subscribers as it would be subsidized entirely
21 by community dollars. Each iamworthit user would be granted free access privileges to the premium content on
22 all sites which are part of the program. Some content may be purchased and/or reusable, other may be entirely
23 site-specific and novel. This model would be particularly appropriate for community sites which are largely
24 member-based (or for example ISP-member-based communities) where much of their value to members is based
25 upon information and other content which it can provide. It is conceivable that all iamworthit-enabled community
26 sites would enable free access to their content by all other iamworthit customers (though it is possible that they
27 may be mutually restricted if members of competitor communities as desired by the community site). If an ISP
28 service is not already provided, a virtual-ISP service could additionally be offered at a substantially reduced price
29 or possible free (depending upon the number of community dollars left over. One could imagine further extending
30 this present network of free content to free content iamworthit subscribers for free access to fee-based television
31 programming or VOD services. In as much as community sites and television channels are becoming different
32 media for delivery of the same information as the number of channels increases, VOD becomes technically
33 enabled, and, most imminently, full motion video can be delivered upon demand over the Web.
- 34 5. Free Access To Subscription and Fee-for-use Information of the Web - In addition to the aforementioned free
35 community site content, it would be possible to further provide free and automatic access to fee-based information
36 on the Web. These costs may be able to be covered by the model across all or most sites depending upon the usage
37 characteristics of its users (e.g. assuming advertiser/community dollars payment to the sites are averaged across
38 users in accordance with the consumption patterns of the average iamworthit user). The identity (pseudonymous)
39 of the user would have to either be disclosed to the site via the proxy or a unique pass code (as required by the site)
40 provided to the user could be automatically entered upon the user accessing the fee for use area requiring the code.
41 A directory (portal) of these fee-based sites would be a useful adjunct to subscribers.
- 42 It would be possible to offer websites the ability to become Internet service providers where the interface to the
43 ISP home page would essentially be heavily branded to that site or portal. Companies like GTE already offer a
44 "Virtual ISP" service in which the content to the ISP home page is unique to the ISP while the network is provided
45 by the virtual ISP service. This model would be particularly compelling for sites which are largely community
46 oriented and have a potentially loyal customer base. Moreover, interestingly, many of these community sites are
47 offering many of the services and capabilities that a full-blown ISP would offer from its home-page, e. g., a portal
48 interface, links to high-quality content, chat/forums, e-commerce, commerce affiliate links, etc.
- 49 6. Bundling iamworhit links with hardware with a PC manufacturer - Many PC manufacturers are now recognizing e-
50 commerce as a very important sales channel. The present model would involve the P.C. manufacturer bundling a
51 link along with a promotion for iamworthit. The promotion would offer the user cash credit for the user. The PC

CONFIDENTIAL

- 1 manufacturer would also receive exclusive advertising rights to target users whose browsing behavior profile
2 qualifies them as a future sales prospect. E.g., instead of cash credit as the profile is identified, the offer could then
3 become modified to offer free hardware or credit to their purchase. Because PCs are highly portable, the
4 advertising targeting techniques described in LEIA could add substantial additional value to advertisers. A similar
5 model could be used for manufacturers of PDAs.
- 6 7. Providing a service to iamworthit partner vendors whereby their users may be automatically matched with each
7 other based upon the similarities of their user profiles. This may be done for purposes of introduction or for
8 introducing users into on-line chat or discussion forums. This service may be provided as a virtual service wherein
9 a menu of different forums and chats are displayed and accessible from each iamworthit member site. (the
10 underlying methodology for which is described in co-pending patent application "Virtual Community Service for
11 System for Customized Electronic Identification of Desirable Objects"). In accordance with this specification, a
12 variation of the service involves the process for identifying individuals who most closely match a given category or
13 target object. For example in the context of the present implementation a category or content, merchandise or a
14 purchasable being specially promoted may be the focal point of a discussion forum or chat room, which is
15 automatically organized by the Virtual Community agent. Accordingly, a portal (or in accordance with the present
16 trend) a site with which a portal interface is integrated utilize the present techniques for generating virtual
17 communities for each category or sub-category of content on the portal or for direct access into a forum or chat
18 room which was automatically created around that particular site (as the target object used as the matching
19 criterion). As described, the user may navigate a hierarchical menu of virtual communities which may be
20 constructed automatically according to the methods described which involves communities assigned to category,
21 sub-category, and association with corresponding sites. Ideally in this scheme the portal is actually a "virtual
22 portal" which may be utilized in providing access to the communities across numerous sites (and/or ISP home-
23 pages). Users may also be navigated (at the individual user level) which along with their pseudonymous user
24 profile data is subject to their data release policies. In a variation of the above schemes, if there is geographical
25 information which is associated and which is released in accordance with the above individuals and/pr
26 communities (e.g., as may be occurring or scheduled to occur in physical space), LEIA may be employed as a
27 primary (or additional) selection criteria for navigating the present information accordingly.
- 28 8. Creation of an iamworthit online multi-store retail site - Instead of selling the community dollars model to on-line
29 vendors, an alternative approach would be to establish a retail presence in a (or potentially multiple) retail niches.
30 The primary business model would be to leverage existing large iamworthit subscriber base (involving the other
31 various types of commercial partners) in order to dedicate a certain percentage of the community dollars (e.g.
32 thirty percent or approximately one hundred fifty dollars per customer) which could only be redeemed at that
33 multi-store retail site (and/or the value of these dollars could be worth more at the retail site). In addition, in this
34 model, the independent advertising initiative of iamworthit would be geared towards community dollar credit of
35 that retail site. It should be noted that, because if other outside competition occurs to the basic iamworthit scheme
36 to a substantial degree there will not be a compelling incentive for users to adopt a more restricted form of value
37 (as retail credits at a particular site), versus accepting the credit from a competitor in the form of cash. Thus this
38 model could provide a viable means for attaining a leading position in one or more on-line retail markets if this
39 competition does not substantially exist.
- 40 9. Advertising in Exchange for Equity - A potentially attractive optional form of value, which could be provided to
41 iamworthit customers involves equity shares in companies which advertise to the user (in lieu of community dollar
42 credit or cash). This scheme is an ideal application for iamworthit in as much as iamworthit customers can be
43 highly targeted and because many Internet-based start-ups are highly niche community oriented (thus iamworthit
44 customers who are interested in the sites can be efficiently identified and targeted). Moreover advertising is
45 typically very expensive which in the absence of accurate targeting may be of questionable value. It should be
46 noted, however, that because the primary objective is to both find viable prospects and to engender an element of
47 loyalty (which the equity model does). This scheme would be the preferred approach to advertising for sites which
48 do not sell on-line where community dollars would be the preferred loyalty engendering scheme. In order for this
49 model to substantially provide its desired advantages of increased advertising exposure to fledgling web based
50 companies, the iamworthit subscriber base would have to be quite substantial.
- 51 10. Loyalty credits for off line retailers - Deliver through the back of sales receipts, kiosks or direct mail or on-line
52 substantial purchase credit to retailers (e.g. grocers) customers, using the aforementioned technique of utilizing a
53 unique URL to identify the vendor and/or promotion from which an iamworthit subscriber originally accessed the

CONFIDENTIAL

- 1 iamworthit subscription site (thus identifying for both user and vendor the appropriate denomination and/or terms
2 of community dollars issued to the user). In the preferred implementation, a loyalty card is used to identify the user
3 thus enabling the community dollars value to be provided to the customer at check-out as straight credit or possibly
4 an enhancement to loyalty credit. The user may also be identified via credit card or alternatively a voucher (or
5 coupon) could be printed from the user's computer or from a kiosk which is typically situated near the entrance to
6 the store and which could be activated upon insertion of a loyalty card credit card (or associated authorization
7 code) and could also be used to disclose the user's community credit balance. A unique identifier for that voucher
8 or coupon is provided and non-tamperability measures are provided such that the user's community dollars account
9 can be appropriately debited upon redemption. Preferably, a pre-determined value is specified on each voucher
10 (which could be predetermined by the service or the user) or alternatively, the total community dollars balance
11 could be specified on the voucher along with the user's name/address and redeemable only upon presentation of
12 valid user ID.
- 13 11. Auto Insurance Application - Co-pending patent application entitled "Applications for a Location Enhanced
14 Information Architecture" describes a location-enhanced framework by which statistical methods are used in order
15 to very efficiently and confidently extrapolate the most relevant attributes in predicting automobile accidents (or
16 the avoidance thereof). The correlations from some of the existing metrics used may be refined using this
17 technique e.g. LEIA is able to accurately determine the number of miles a user drives per week while the user will
18 often lie about this, thus the basic model may be refined and more accurate information may be provided on a per-
19 user basis. The scheme also enables completely new metrics to also be identified and utilized as well which may
20 correlate the attributes location with time. It is conceivable that if a user provides access to this location-enhanced
21 information by an insurer, that the insurer could in turn offer premiums, discounts or deliver credit to the user
22 which could be added to monetary credit the user receives for personal information from iamworthit, for example,
23 an iamworthit implementation which uses LEIA to profile and target users with ads by their location (e.g. while
24 riding in an automobile).

25 15. Secure Data Interchange: General Examples

26 15.1 Assessing the Value of Data

- 27 • Plug together sets of data, and measure predictive accuracy.

28 15.2 Matching Data Across Vendors

- 29 • find patterns in common pseudonyms, denoting common areas of interest.
- 30 • use catalogues of order codes and item description to find similarities across data sets

31 15.3 Targeted Recommendations

- 32 • describe CDNow style project : similar customers defined by nearest neighbor and clusters; these are used to
33 create recommendations for an individual customer.

34 15.4 Leveraging Portal Data

- 35 • data from portal to leverage data needs for ISP

36 15.5 Automated Generation of Customized Web Pages

CONFIDENTIAL

1 Analyze customers for broad preferences in choice of web pages visited (corporate, Star Trek fan, etc.). This defines
2 the initial look and feel for the page that greets them at their portal (a teen might enjoy lots of bright colors and sound
3 clips, an investor would prefer a more staid design); different "skins" could be created to match the major categories of
4 customers, and would designate both the graphical design and modules available on the page (e.g., a working stock-
5 ticker for an investor, a real-time weather map for a jogger).

6 The web pages and information most frequently accessed by a customer would be given priority, and a hierarchy of
7 usage could be developed. Since stock prices are of the highest importance to an investor, a ticker reflecting his
8 portfolio value would stream across the top of the page. However, although he enjoys spending his profits on vacations
9 and automobiles, these are only of secondary interest to him (as revealed by his on-line behavior), and so are relegated
10 to a sub-menu on his web-page. As his usage changes, the priority level assigned to the modules would change as well,
11 so that when a jogger purchases a treadmill for indoor running, his weather reports won't dominate the top-level screen.

12 Small children could have simplified browsers, with extra-big buttons and access to pages pre-screened by a "web-
13 nanny" service.

14 SDI would be used in the initial phases to group customers into general categories based on their patterns of their web
15 surfing, and would be used in later phases to adjust the content and style of their portal home-pages (based on what
16 similar customers seem to be enjoying).

17 <<Mention fact that broadvision is trying to do some of these things .broadvision-like: (develop model for automated
18 generation of customized web pages)

19 15.6 Analyzing Affinities

- 20 • suppose a vendor has a list of customers, and knows to some degree what web pages they visited after leaving
21 vendor site. A large collection of customers taken from an ISP will contain their web-surfing behavior. Cluster web
22 sites and cluster customers, finding cluster-to-cluster interactions.
- 23 • use this information to classify vendor's customers; gives vendor an edge in knowing customers' tastes.

24 15.7 Personalization of information.

25 Personalizing information on-the-fly requires that a vendor has a data model, for example that clusters its current user-
26 base according to what they are likely to be interested in. Notice that it is only with more information from outside of
27 their domain that users can be clustered with respect to their product space. Technology One: Need to be able to update
28 vendor's data models, re-cluster users, on the basis of wider information about the users, without revealing that
29 information. i.e. just provide the results of the analysis to the vendor, on his current user-base. This will not violate
30 privacy, and allow rapid personalization. Technology Two: Need rapid profiling of a new user, in a way that does not
31 reveal personal information to the vendor, i.e. can we provide an algorithm that the vendor can run on
32 cryptographically secured profiles that accompany a new user to enable that user to be categorized. Alternatively, the
33 vendor would need to request a categorization for the new user from SDI.

34 15.8 Enable user's to find appropriate information/products.

35 This could be more contentious with vendors, for example if we never recommend particular services.

36 15.9 Ad networks.

37 We need to allow ad networks to show ads relevant to the current user on the page. Approach: each ad network has a
38 set of ads that it is currently running on a page. Can off-line develop a decision tree to decide how to assign a user to an
39 ad, given a profile. Again, we need the user to come with the profile, and the ad network to have a secure evaluation
40 technique that can run from profile, get answer, without getting any data.

CONFIDENTIAL

1 Virtual ad networks. Track users across multiple domains. Leverage the user database. Fine level control over ads that
2 the user sees. The trusted secure data interchange can operate as an "ad network", allowing for the placement of well-
3 focused banner ads to market goods that are relevant to users of a particular content site. Electronic banner ads provide
4 the potential for one-to-one marketing, when the advertising agency has information about the user that has just hit a
5 site, together with information about what the user is doing local to a site. For example a car manufacturer is able to
6 place a focused advertisement to a user that has just performed a search for new cars in a search engine, to a user that is
7 known to have a large family and a high disposable income.

8 There are two possible business models. Firstly, an Internet content provider could purchase access to information
9 placed by vendors and users within the Secure Data Interchange database. This information may be "rented" for a
10 period of time, and then whenever a user visits the site of the content provider (possibly through the pseudonymous
11 proxy server), the provider can query the data interchange for information about the user. The Internet content provider
12 sells well-directed advertisements to vendors. Secondly, the data interchange could sell or rent data to an advertising
13 agency directly, providing information in real-time to enable the advertising agency to provide more focus in its banner
14 ads for its clients. "Per-transaction" pricing is a very powerful pricing model that is enabled with on-line banner ads. It
15 is simple to monitor the number of click-throughs that are received at a particular banner, in response to an
16 advertisement. In the off-line world pricing must be based on the number of impressions, or worst still, the number of
17 mailings sent and it is more critical to understand the expected value of a campaign up front.

18 The proxy server could also act as an "ad network" itself, and sell focused advertisements for vendors, and purchase
19 ad-space on the sites of content providers. The on-line domain provides this unique opportunity for quick
20 experimentation with advertising strategies in order to get feedback on the likely utility of untested approaches. The
21 system can use a hierarchical cluster tree to identify the most revealing items in a dynamically responsive fashion such
22 that the profiles of all of the selections can be generated with the most minimal amount of interactions with the user
23 (see "Rapid Profiling" section in issued patent entitled "System & Method for Customized Electronic Identification of
24 Desirable Objects). Thus a more robust statistical model across multiple vendors is established as a result of the user's
25 click through response of these intelligently selected virtual banners as well as other pages which are subsequently
26 navigated through once the remote site is accessed via the banner.

27 In the preferred approach rapid profiling not only dynamically identifies and presents items which are most revealing of
28 the other items in the collection, it also selects the users whose profiles suggest the greatest familiarity with these items
29 (i.e., potentially correlated items). Furthermore, if the system's objective is to find new users or users who may be
30 interested in the present vendor's other products, products for which little is known, then it will match users who are
31 least familiar with exemplar items. The idea is to reveal the most significant data about the user profile with respect to
32 the present collection of items of interest. Finally, rapid profiling can use direct explicit queries to determine interest
33 on an item(s) or to collect demographic data on a user.

34 The target object profiles of advertisements on the ad server are matched against the user profile in order to
35 automatically present the most relevant recommendation(s). Typically, the client-side proxy requires the host-level
36 proxy to disclose the target object profiles of the products/services sold by the vendor. This data is stored as meta-tags
37 in XML form and is encrypted. This data can be very useful to the user in navigation, filtering and search activities in
38 the future or in a variation the ISP - level proxy a party (a neutral server) could store these target object profiles and
39 selectively disclose relevant pieces of them (e.g. genre cross-correlations) to vendors, which are
40 considered according to the disclosing vendor's data disclosure policy acceptable to receive this data. These profiles
41 are not accessible to the client-level proxy but may be disclosed only if there are restrictions within the vendor's data
42 disclosure policy.

43 In another variation, if the data to be disclosed to the vendor is acceptable to the original vendor but she/he is untrusting
44 of the vendor, the data is received by the host-level proxy (another neutral third party) instead of the vendor, thus
45 providing the disclosing vendor with an additional level of security, assurance about the use of his/her data while
46 enabling the users of such a site to access all of the merchandise or content in a completely personalized fashion. Thus
47 these XML tags are stored in association with, but on a separate server from the actual HTML pages stored on the
48 vendor's site. Additionally, these profiles are constantly updated by user profile data conveyed to the host-level server
49 which operates in distributed fashion.

50 15.10 Personalization of links

CONFIDENTIAL

1 Adding value via data collected about this user and other users. Not just surfing data, but wider data-purchasing
2 behavior etc. On-line commerce enables the dynamic personalization of retail for each user. A virtual shop floor can be
3 arranged to match the predicted preferences of each user. A commercial Internet site can leverage two types of
4 information in order to personalize its product. For a new user, that has never before visited the site, it is very
5 advantageous for the site to already know about the preferences of that user in order to personalize the goods and
6 services that it offers. The information provided at the secure data interchange, and gathered from the transactions of a
7 user with another vendor, is vital for this type of personalization to first-time users. The second type of information that
8 an Internet site can leverage is information that it has collected from previous interactions with the user, information
9 that is collected locally to the site.

10 15.11 Hospital database

11 In one data application of the client level proxy server, the user profile includes medical data which is obtained from
12 medical records. (such as from hospitals or physician's medical records or potentially that of a health insurer).
13 Typically, various physician's offices and hospitals which a patient (hereinafter "user") has visited over the years
14 contains separate portions of a user's overall medical history, thus these various sources may be combined upon the
15 user's request by downloading this data to the client-level proxy (or preferably, the user enters into a contract with
16 those organizations in which all medical data and updates thereof are downloaded by the organization and/or an
17 "agent" to the organization which transmits a request which is digitally signed by the user at the client-level proxy
18 server. The origin of the request (the user) is authenticated and may be processed by a human or another agent located
19 at the organization's host computer.

20 Because of the highly sensitive nature of medical data, there are potential user privacy advantages in using randomized
21 aggregates. For example, a user's age, medical history of specific relatives (particulars of which could be more
22 generalized) genetic data, numeric values associated with various medical tests, results for which are a numeric value.
23 This data may be of relevance to pharmaceutical companies, alternative medicine vendor and clinics insurance
24 companies hospitals physicians, clinics and home health care providers, the latter three of which may wish to advertise
25 to patient prospects and extend their medical practices. The privacy architecture herein provided is a critical
26 component for enabling access to user data by these commercial entities.

27 15.12 Smart Web browsing

28 Definitions: exemplar - the profile of target object or (as pertinent to following description), user profile which is
29 "most like" the profile of the cluster to which it belong, perhaps a median metric.

30 The Platform for Privacy Preferences (P3P) provides for the ability to utilize XML meta-tags for purposes of
31 annotating Web pages with comments from previous visitors to those pages. It further suggests that users with an
32 affinity or otherwise associated with a particular category (e.g. of expressed interest) may both identify themselves as
33 well as select a category (e.g. which they are associated with) and observe the annotations of other users associated
34 with that category. One of the divisional applications of the parent case "System for Customized Electronic
35 Identification of Desirable Objects" relating to the automatic creation of virtual communities suggests that users may
36 be automatically assigned to particular communities (e.g. chat groups, forums etc.) which may be either automatically
37 generated and labeled or constructed and labeled manually.

38 Additionally, users could further be allowed to rate the annotations to the pages, which could be a means by which
39 qualifications of users to provide annotations could be measured by content domain (i.e. cluster). Their comments (and
40 particularly future comments) could then receive a priority position in the annotated comments available. Future
41 comments from users with a poor rating history for a particular content cluster may be deleted. A persistent interface
42 feature on the tool bar or side bar may provide for annotations to also be accessed by users selecting certain profile
43 features of users as they browse from page to page. E.g. identify the comments of a news article about abortion by
44 users who are self identified as advocates of the Women's Rights Movement, ultra conservative senior citizens, teen
45 women or those with a strong interest in alternative medicine or the Catholic Church or identify the annotations relating
46 to Ford Motor Company by general Motors (though in the case of competitor annotations the above description
47 suggests means protective safeguards for the recipient of the annotation).

CONFIDENTIAL

1 The parent case further suggests that users may actively provide ratings in a completely privacy protected manner
2 according to various criteria of pages they browse. A reasonable extension would further include being able to observe
3 how these ratings relate in accordance with various user profiles (or groups which users identify themselves in
4 association with). E.g. overall quality, aesthetic appeal, interest in the content, value (if it is a purchasable) etc. Each of
5 these criteria's overall ratings would vary in accordance with the particular user groups which is selected by the user.
6 Accordingly users may submit as a query a user profile or user cluster profile a page rating criteria a combination
7 thereof and receive a listing of the pages of highest relevance to the search criteria. These may include, for example,
8 the exemplar user profiles of users who previously most visited the site or the exemplar user profiles of the most
9 predominant clusters of users who previously visited the site. In addition, the page rating profiles may be also
10 displayed as well as related links which are determined to be most relevant to that collective group of users e.g. as
11 statistically estimated from the referral logs or explicitly identified as book marks by those users as being of particular
12 relevance or similarity to the present page.

13 As suggested and described in the parent issued patent application, these clusters may be accessed through querying
14 (search) filtering (where relevant pages and/or annotate to relevant page are "pushed" to the user as they appear on the
15 Web), and browsing which includes navigating a hierarchical menu of users who are classified according to their
16 passive behavior patterns and/or ratings which have been actively submitted as well as automatic hyperlinking (or
17 automatically linking the presently viewed page to its "nearest neighbor"). Based upon the above criteria as selected by
18 the user, ratings and annotations may be viewed in addition to these techniques are preferably deployed in the context
19 of the navigational techniques as taught in the parent case.

20 The above description also describes the use of a hierarchical menu through which groups of users may be identified by
21 their profile features (wherein a profile feature could even be a rating criteria itself of for e.g. an opinion via a site
22 survey). These features could also include a user selected rating criteria by a certain type or group of user or users
23 sharing rating similarities. These features could in turn be used to either selectively filter-out content which falls
24 outside of that criteria as the user navigates the information by a variety of means (e.g. it could be used to selectively
25 filter items or even categories in the hierarchical menu) or identify if/when pages are otherwise encountered where
26 these user rating features are present (or metric features are predominant), thus displaying this user statistical
27 information in conjunction with the ratings statistics and/or associated annotations if desired. (NOTE: CLARITY ??)

28 Other criteria for observing ratings and annotations may be in accordance with those submitted by organizations. As
29 above described, recall that the user may also use one or more of the organizational approval criteria to also bias or
30 completely filter selections as the user proceeds to navigate the Web. These endorsements may apply to individuals
31 which may either provide annotations, ratings, news group postings, or editorial content.

32 Vendors delivering targeted on-line advertisements may also be interested in the above information, in particular the
33 profiles of the most exemplary groups of users and their affinity ("similarity") toward particular ads, purchased
34 products or product or content categories—as measured against user attributes or exemplar user profiles of the most
35 relevant clusters of users based upon passive behavior and/or active ratings. This information may, of course, may be
36 displayed (or selectively displayed) with the pages for the product (or service) as meta-data in accordance with the
37 vendor's wishes. For vendors, these techniques may of course be applied exclusively within the scope of the web-
38 sites of the vendors concerned (and/or additionally to users). For example, the site-specific page view correlations
39 (including time spent viewing each page) in accordance with the dominant clusters, Exemplar user profiles and
40 attributes of those users are certainly of interest to vendors to which those sites belong as well as affiliate sites on
41 which their advertisements and/or syndicated products are advertised and sold remotely.

42 15.13 Location Enhanced Delivery System Architecture (LEIA) Enhanced SDI

43 For an exemplary application of the Secure Data Interchange technology, consider its extension to co-pending patent,
44 entitled "Location Enhanced Delivery System Architecture" (LEIA), we teach a method for matching information
45 providers and information recipients that utilizes location information, in addition to static and dynamic profiling
46 information. The method customizes the information that is displayed on a private or public information device to the
47 real audience in the vicinity of the device, instead of a predicted audience. LEIA collects an extremely detailed and
48 comprehensive information set about the daily activities of a user, enabling enhancement of the user profile with

CONFIDENTIAL

location information and temporal activity patterns. The co-pending LEIA patent suggests appropriate application environments, for example in a smart home, an office, on a mobile shopping device, and in an automobile. A LEIA-based system stores personal information on users.

We can extend LEIA by incorporation with the Secure Data Interchange system that we teach in this patent. SDI enables the user to receive the benefits of powerful and well-directed information, but within a system that respects his/her privacy requirements. The interchange acts as a secure data warehouse for users and information providers, enabling information providers to target users without revealing private information to the providers directly.

Co-pending patent, "Location Enhanced Information Delivery System Architecture" (LEIA) customizes information that is displayed to an information recipient based on object profiles and physical location of users. Presents the information most relevant to the REAL audience, not a predicted audience.

One application includes "Smart Home Intelligence", where methods are disclosed by which users' real-time behavior may be profiled through their movement throughout their home, and specific interactions with the various network enabled appliances throughout the home. Other inputs may include the user's speech patterns (using voice recognition and text analysis). It could for example, note the user's speech content patterns in real-time. Such information provides invaluable clues as to the user's present activities, mood and interest state and may be processed by the presently described algorithms tuned with location/time features typically using the assistance of human data analyst to identify the key features and correlations. (This information may also provide enhanced information pertinent to the user's general, static preferences as well).

Other extensions of this scheme are also considered e.g. within the context of the user's office, or automobile and pedestrian activities. This application may thus extend the usefulness of the iambworthit model to advertisers in being able to target users through the presently anticipated on-line media as well as networked appliances and in either case, based upon the relevant context of users' present activities and behavior (and from this potentially their inferred moods or mental states) within their homes and elsewhere.

15.14 Extended Example: Vacation Packages

A vacation package organizer decides to begin a large-scale marketing campaign to target those people who would be the most interested in joining a new Caribbean Cruise. Although the vendor has a database of current customers, it is interested both in increasing the number and suitability of its potential leads.

Interfacing with the secure data interchange with which it is a member, the organizer identifies several possible sources of supplemental data: a LEIA-based travel discussion group, an on-line bookstore, and a Caribbean restaurant. These are found both by browsing through the interchange's internal list of members, and by using SDI-based data analysis tools, used within the interchange to automatically identify entities sharing common characteristics.

The package organizer then contacts each of these entities through the interchange, and negotiates different data-sharing deals: the travel discussion group is willing to exchange full information for a large travel discount, the on-line book store is willing to reveal the pseudonyms of users who have bought travel books in exchange for a per-sale commission, and the restaurant is willing to sell its entire database for a flat fee (and will provide an aggregated data set as a sample).

The vacation package organizer now chooses fairly basic data-mining algorithms to identify the individuals with the greatest potential interest in a Caribbean vacation; however, the organizer does splurge on a new neural network approach developed by a small software company. On a per-sale commission, the Software Company is willing to loan the vacation package organizer use of its data mining code.

First, the organizer decides which data sets to use. The initial results on the restaurant's aggregated data aren't so good (its customers turn out to not be very affluent), so the organizer declines the purchase of the full data set. However, it does agree to the conditions asked by the travel discussion group and the on-line bookstore.

The data provided by the discussion group and on-line bookstore, being in a common format, are moved in a secure fashion to the interchange's processing area, and are acted upon by the data mining tools, which are also in a compatible format. As per the agreement, the interchange forwards discounted Caribbean cruise offers to the members of the discussion group, and forwards standard promotions to targeted individuals in the book store's customer list. A few of these individuals respond favorably; these electronic transfers of money are passed back through the

CONFIDENTIAL

- 1 interchange, which slices off a commission for the book store before passing the accepted offers back to the tour
2 organizer, who learns the identities of the customers and can now count them as part of its database.
- 3 This protocol specification could even be digitally signed by the "owner" of the data as proof of ownership of the data
4 and its associated restrictions by the owner, i.e., effectively a "digital deed" which is both legal and untamperable by
5 any other party and thus acts as a legally binding proof of ownership and terms/conditions dictating how that data can
6 be used.

7 16. Future Directions

8 (i) Digital Ownership Deeds.

- 9 A digital-deed is a digitally signed credential issued by a trusted witness of a transaction. It can be useful to have a
10 secure record of a large ticket item that is bought or sold in an on-line transaction. It is convenient to have this verified
11 and maintained, automatically. Provide increased confidence for members of SDI, particularly those which may wish
12 to have a record or "back-up" verification of large transactions e.g. within the business to business implementation of
13 SDI or for real-estate transactions, investments in small entities, etc. Additional data may further be associated as part
14 of the digital deed credential, including any restrictions, stipulations, (including time-based), warranties, insurance, etc.
15 pertaining to that item. Credentials cannot be tampered with by some other party, secured using cryptographic
16 techniques.

- 17 In the domain of "information goods", SDI can automatically add an electronic "water mark" to digital information
18 (e.g. a video, a CD, or a book), that associates a license for use, and terms-of-use. In this way (because the water mark
19 cannot be removed, but does not affect the use of the data), the legality of selling/duplicating pure information goods
20 can be verified by inspection of the water mark. Metadata can be used to assign licenses for intellectual property, along
21 with data regarding terms of licenses, assignments, filing and expiration dates, etc. The seller can impose, within this
22 untamperable credential, certain terms and conditions of who the buyer may be.

- 23 Credentials can be used to restrict the ability of minors to purchase pornographic, or other unsuitable material; and can
24 also be tagged to products that are sold, so that such material cannot be exchanged within SDI without the correct
25 credentials.

- 26 Credentials can also be used to create a "resale market" for potentially any and all items a user purchases (which has a
27 potential resale value). The user that buys the good can be an "advertiser" in a resale market. The user will want to
28 prescribe controls over the personal information made available to interested parties in the resale market. A user
29 becomes a potential seller upon having made a purchase for an item (which is documented and recorded with the digital
30 deed).

- 31 She/he is automatically asked upon transaction (by software on the client-level or ISP-level proxy) whether she wishes
32 to have his/her item listed as a potential purchasable. If no, she/he is asked if/when at a later time she/he may change
33 his/her mind. A typical price range for that type of item is presented and a question is asked about an approximate price
34 range she/he would be interested in selling for (though a stated range is optional since the preferred variation involves
35 the use of a bidding scheme). The user may wish to control identifying information about themselves, i.e. they may
36 wish to remain anonymous or pseudonymous. The current seller can however be provided with assurances about the
37 user, before entering into a commercial relationship.

- 38 Regarding prospective future buyers, so long as the future transaction occurs on line (or conceivably even if it occurs
39 off-line), a user credential can be a required stipulation to ownership of a given product. For example, some vendors
40 may wish to impose a restriction which states that any product which they sell may not be sold to a competitor (based
41 on stated characteristics and/or an explicit list of competitive vendors), or this will prevent an on-line black market for
42 age-restricted goods. Such goods, e.g. information goods, can be protected at source—and water-marked to prevent
43 unidentifiable duplication. Similarly, users (e.g. owners of kittens) may have certain personal interests whereby items
44 of personal or sentimental value which are sold should be owned exclusively by certain types of users.

CONFIDENTIAL

1 Another example may be the sale of corporate assets, whereby the ability to impose certain restrictions on these assets
2 may provide negotiating leverage at the present time for the seller. The seller may state that if a presently submitted
3 offer under the stated terms is not accepted by X (future) date, the offer will be cancelled. The seller can credibly
4 commit to this negotiation strategy, via software credentials. SDI acts as a trusted mediator.

5 A seller could also make a believable threat to place a penalty on all buyers, with except for on some commercial
6 entities, through irrefutable statements. The excused entity could be a competitor to the prospective buyer. The seller
7 can also commit to a time limit, such that the seller cannot reverse the time limit. In general, the ability to make
8 irrefutable claims improves the efficiency of negotiation.

9 Within the present system SDI provides the framework by which appropriate buyers and sellers may be matched
10 together. It also enables a methodology by which the buyer interests are protected through the use of matching of
11 sellers offers to competitive vendors (using iamworthit).

12 (II) Monitor Privacy Violations

13 We can also monitor sites to determine privacy violations. A human analyst can review the stated privacy policies of a
14 Website, and monitor information that is used about an individual, from the personalization that a user receives. If the
15 practices are consistent with stated policies then a credential can be issued. Otherwise, it may be possible to passively
16 observe the degree (e.g. frequency and nature) of the violation, and what is the particular privacy violation. This can be
17 done in combination with any history and details of privacy litigation (e.g. damages). This information could be very
18 useful to a privacy insurer in determining at any given time, what sites are insurable and the associated risks of insuring
19 them.

20 This information may also be used to indicate to iamworthit users (e.g. statically via a "black list" or dynamically
21 while browsing) what sites violate user privacy, what have been the nature and (potential) extent of violations to other
22 users as delivered via the client-level or ISP-level proxy. The subscribers may choose to adjust their data disclosure
23 policy settings to assume (for example) single site pseudonymity or complete anonymity whenever visiting those sites
24 and only certain profile information (or none) may wish to be released accordingly. This information release benefits
25 also iamworthit in as much as in the preferred implementation, it automatically provides its users with privacy
26 insurance, paid from a share of ad revenues (unless the user electively opts out).

27 We can also provide reports to users, including specific information relating to their own interactions and transactions
28 with a vendor. The information may include pageview statistics, time spent viewing the site or its pages by visitor or
29 collectively, transactions and amounts transacted. In a variation, the service may provide user specific information
30 regarding perks, benefits and bonuses (e.g. community dollars/discounts) which the user is entitled to.

31 ** DP. Technical disclosure for above?

32 (III) Vendor-centric Tools

33 Finally, metrics are useful within electronic commerce—for example, how many hits does an advertisement receive,
34 what is the profile of users that hit the advertisement, how well does targeted advertising perform, etc. This information
35 is available within SDI, through client-level proxy monitoring, and can be released or aggregated, according to user
36 data-release policies.

37 There are a variety of different auditing, validation and reporting services which can conceivably be performed and,
38 furthermore, this data can be analyzed and converted into digital credentials which may be useful to sites. Other issued
39 directly to the site for the benefit of validating the associated information to visitors to the site (these credentials could
40 alternatively be displayed virtually via the ISP-level proxy).

41 Auditing and validating-click-through rates of various on-line vendors and advertising services. Data collected by
42 iamworthit at the ISP-level proxy server and by the vendor-centric SDI service at the host-level proxy may provide a
43 form of a reporting service in which SDI could be used to perform on-line usage and transaction measurements across a
44 variety of sites (including competitive sites personalization enabled and targeted ad delivery sites etc.)

45 Existing reporting tools are typically implemented on a vendor's own site, however, the present scheme which monitors
46 the user at the click-level proxy is valuable for purposes of suggesting correlations of visitation and behavior relating to

CONFIDENTIAL

1 sites external to the vendor with those simply providing raw aggregates statistics regarding click-through clusters and
2 associated traffic volume on that particular site sometimes as a fraction of time.

3 The primary goal of SDI is to protect the privacy of users, while providing incentives for users to provide information
4 (within a controlled environment), and allowing users to reveal personalization within a controlled environment.
5 However, it is conceivable that vendors might try to break the system, and make links between users across different
6 vendors. Relevant information could include: the timing of access to and from sites, the time between clicks, typing
7 cadence, and cursor movement. The system of SDI disables the Netscape cookie mechanism, removes referral tags
8 from HTTP messages, and anonymizes routing information — however if any one of these systems is not in place
9 (perhaps a user is not subscribed to SDI) then vendors can clearly use all of these to gain additional tracking accuracy.

10 The Vendor-centric version of SDI (unlike the user-centric version i.e. iamworthit) may wish to make optimal use of
11 information which is collected from an unidentified visitor's click stream in order to make inferred "guesses" about
12 who the user might actually be (e.g. a user who is unidentified by cookies, digital certificates, their client-level proxy or
13 their customer account or credit card number). The vendor-centric version of SDI may utilize the data it collects from
14 users across multiple vendors in order to make certain inferences about the identity of that individual.

15 A user's overall profile (available to all vendors) may reveal clues about the identity of user, if it is not sufficiently
16 randomized.

17 ** With vendor cooperation, referral across sites can be used—in conjunction with timing clues. This can be very
18 valuable information, in conjunction with general profile information, similar typing, click rates, interests etc. i.e.
19 vendor cooperation can recreate some of the information that is lost when the HTTP referral log is carefully "washed"
20 via the SDI proxy server.

21 A robust user profile, across multiple vendors, is very valuable. Such a profile can be very revealing of a user's
22 identity, allowing a more complete profile of the user's interest. Many of the same techniques described in SDI to
23 extrapolate user preferences (including data mining, collaborative filtering and text analysis) may uncover unique
24 identifying features.

25 Additionally, the fact that this information may be correlated with other unique identifying features, such as dynamic
26 manual features of click stream, (time between clicks), cursor movements and typing cadence may enable the
27 construction of a rather complete picture of the user and his/her identity. (e.g. using neural nets..)

28 A user could also be presented with questions to answer, to establish identity—combining a certain amount of real data
29 from a user's SDI proxy with the vendor's profiles of users.

30 This can help to positively identify a user across sites.

31 Multiple cooperating sites may use the present techniques to predict, both when a user is likely to have returned to
32 his/her site, and when a user is likely to have accessed another vendor site. This can allow a vendor to update a user's
33 profile (with cooperation with another vendor), on the basis of what a user does at the subsequent site. Profiles can be
34 updated statistically to allow for uncertainty.

35 Cookies cannot be leveraged under the standard SDI model, because they are disabled.

36 Theoretically, all vendors could, cooperate in the present initiative (to identify and comprehensively profile the user
37 without his/her consent) while not authorized SDI to do anything further with that data (for data sharing purposes). If
38 each vendor logs the path of users (under pseudonyms) as they pass from site to site, it could be possible to use data-
39 mining techniques (with timing patterns) to make statistical predictions about user pseudonym portfolios. However,
40 data is likely to be noisy, and the amount of data that is necessary is huge, especially when many users may hit a
41 vendor's site simultaneously.

42 17. References

43 [APPEL] "A P3P Preference Exchange Language (APPEL)." World Wide Web Consortium Working Draft (Work in
44 progress), 1998. URL: <http://www.w3.org/TR/WD-P3P-preferences>

CONFIDENTIAL

- 1 [CKR97] Connolly, D., R. Khare, and A. Rifkin. "The Evolution of Web Documents: The Ascent of XML." *World*
2 *Wide Web Journal* 2 (Fall 1997): 119-128.
- 3 [IRDF] Lassila, O. "Introduction to RDF Metadata." W3C NOTE 1997-11-13. URL:[http://www.w3.org/TR/NOTE-](http://www.w3.org/TR/NOTE-rdf-simple-intro)
4 [rdf-simple-intro](http://www.w3.org/TR/NOTE-rdf-simple-intro)
- 5 [KRR97] Khare, R., Rohit, and A. Rifkin. "X Marks the Spot: extensible Markup Language opens the door to a
6 motherlode of automated Web applications." *IEEE Internet Computing* 1 (July/August 1997): 78-87
- 7 [KR97] Khare, R., and A. Rifkin. "Capturing the State of Distributed Systems with XML." *World Wide Web Journal* 2
8 (Fall 1997):207-218.
- 9 [MCF] "Meta Content Framework Using XML." Submitted to W3C, June 1997.
10 URL: <http://www.w3.org/TIUNOTE-MCF-XML-970624>
- 11 [MCFXML] Guha, R.V., and T. Bray. "Meta Content Framework using XML." Submitted to W3C, June 1997. URL:
12 <http://www.w3.org/TR/NOTE-MCF-XML/>
- 13 [Mil98] Miller, E. "An Introduction to the Resource Description Framework." *D-Lib Magazine* May 1998.
- 14 [NS] Bray, T., D. Hollander, and A. Layman. "Name Spaces in XML." W3C Note, January 1998.
15 URL:<http://www.w3.org/TR/1998/NOTE-xml-names-0119>
- 16 [OPS] "Proposal for an Open Profiling Standard." Submitted to W3C, June 1997. URL:[http://www.w3.org/TR/NOTE-](http://www.w3.org/TR/NOTE-OPS-FrameWork.html)
17 [OPS-FrameWork.html](http://www.w3.org/TR/NOTE-OPS-FrameWork.html)
- 18 [PRIVACY] "Privacy and Profiling on the Web." Submitted to W3C, June 1997.
19 URL: <http://www.w3.org/TR/NOTE-Web-privacy.html>
- 20 [SCHEMA] The W3C RDF Schema Working Group,
21 URL:<http://www.w3.org/TR/WE-RDF-Schema/>
- 22 [SGML] *information Processing—Text and Office Systems—Standard Generalized Markup*
23 *Language (SGML)*, International Organization for Standardization, Rdf No. ISO 8879:1986, 1986.
- 24 [SPEC] "Resource Description Framework (RDF) Model and Syntax."
25 URL:<http://www.w3.org/RDF/Group/WD-rdf-syntax/>
- 26 [URI1] *RFC 1738: Uniform Resource Locators (URL)*. IETF RFC1738 IETF (Internet Engineering Task Force, ed. T.
27 Berners-Lee, L. Masinter, and M. McCahill. 1994.
- 28 [URI2] *RFC 1808: Relative Uniform Resource Locators*. IETF RFC1808 IETF (Internet Engineering Task Force). ed.
29 R. Fielding, 1995.
- 30 [VC] *vCard Home Page*, URL:<http://www.imc.org/pdi>
- 31 [W3C] *The World Wide Web Consortium Home Page*, URL:<http://www.w3.org/>
- 32 [XML1] Bray, T., J. Paoli, and C. M. Sperberg-McQueen. "Extensible Markup Language (XML): Part I. Syntax",
33 World Wide Web Consortium Working Draft (Work in progress), August 1997.
34 URL:<http://www.w3.org/TR/WD-xml-lang.html>
- 35 [XML2] The W3C XML Extensible Markup Language Working Group Home Page,
36 URL:<http://www.w3.org/XML/>

85

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.